NASA / CR-1998-208455

A Method for Making Cross-Comparable Estimates of the Benefits of Decision Support Technologies for Air Traffic Management

David Lee, Dou Long, Mel Etheridge, Joana Plugge, Jesse Johnson, and Peter Kostiuk
Logistics Management Institute, McLean, Virginia

July 1998

The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counter-part or peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that help round out the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at ***http://www.sti.nasa.gov***

- Email your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA Access Help Desk at (301) 621-0134

- Phone the NASA Access Help Desk at (301) 621-0390

- Write to:
  NASA Access Help Desk
  NASA Center for AeroSpace Information
  7121 Standard Drive
  Hanover, MD 21076-1320

NASA / CR-1998-208455

A Method for Making Cross-Comparable Estimates of the Benefits of Decision Support Technologies for Air Traffic Management

David Lee, Dou Long, Mel Etheridge, Joana Plugge, Jesse Johnson, and Peter Kostiuk
Logistics Management Institute, McLean, Virginia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS2-14361

July 1998

Available from the following:

NASA Center for AeroSpace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

Abstract

We present a general method for making cross-comparable estimates of the benefits of NASA-developed decision support technologies for air traffic management, and apply a specific implementation of the method to estimate benefits of three decision support tools (DSTs) under development in NASA's Advanced Air Transportation Technologies Program: Active Final Approach Spacing Tool (A-FAST), Expedite Departure Path (EDP), and Conflict Probe and Trial Planning Tool (CPTP). We also review data on the present operation of the national airspace system (NAS) to identify opportunities for DSTs to reduce delays and inefficiencies.

Contents

FIGURES

TABLES

Chapter 1

Introduction

This report's chief objective is to present a method for making cross-comparable estimates of the benefits of decision support technologies (DSTs) for air traffic management. As described in Chapter 3, the method uses an economic model to integrate estimates of DSTs' effects on various aspects of the national airspace system (NAS) into a single dollar benefit figure. The underlying estimates of the DSTs' impacts come from a suite of models, including models of aircraft performance, airports, Terminal Radar Approach Control (TRACON), and enroute Air Traffic Management (ATM) sectors, future demands for ATM services, and weather models.

The method itself is a general one. Many different models and data sets can be used as its component parts. In this report, we instantiated the method with a specific set of components. As described in Chapter 3, these are the Base of Aircraft Data/Flight Segment Cost Model (BADA/FSCM) of aircraft performance; the LMINET queuing network model of the national airspace system, which includes weather-responsive models of airport capacity for the network's 64 airports, as well as queuing models of TRACON and enroute sectors; the Federal Aeronautics Administration's (FAA) forecasts for both ATM demands and capacity-increasing projects; and weather data from the National Climatic Data Center. With the chosen components, we modeled three DSTs—Active Final Approach Spacing Tool (A-FAST), Expedite Departure Path (EDP), and Conflict Probe-Trial Planning Tool—and we generated the cross-comparable benefits estimates of Chapter 4.

A secondary purpose of the report is to explore available data on present operations of the NAS to identify opportunities for NASA-developed DSTs. Given in Chapter 2, that material includes surveys of data on airport groundside congestion, on as-flown routes from the FAA's Enhanced Traffic Management System (ETMS), on as-planned routes from the United Airlines' Air Operations Center, and on terminal-area airside congestion for both arriving and departing flights (from ETMS data). Chapter 2 also includes suggestions for DSTs made by a controller at the FAA's Command Center.

Chapter 2
Causes of Delay in the National Airspace System

This chapter surveys data about the present operation of the national airspace system to identify causes of inefficiency and delay that NASA-developed DSTs might address.

After a review of airport groundside congestion, we consider the efficiencies of as-flown routes and the routes planned by a major carrier’s Air Operations Center. Then, we identify costs of the effects of terminal area airside congestion on departures and arrivals. Finally, we report a suggestion from an FAA controller for a DST to improve sequencing flights into, and out of, holding patterns.

It is important to note that the delays and inefficiencies that we consider in this chapter are those experienced by operating flights, between gate departure and gate arrival. The very significant delays that are taken in ground holds do not appear here. These are treated in the benefits estimates of Chapter 4.

AIRPORT GROUNDSIDE CONGESTION

To explore inefficiencies and delays from this cause, we examined taxi-out and taxi-in delays as reported in the FAA’s Performance Analysis Monitoring System (PAMS), which is based on Airline Service and Quality Performance (ASQP) data. Unfortunately, the “taxi-out delay” data confound delays in queues for departure runway service with delays actually caused by inadequate taxiways, and the “taxi-in delay” data confound real taxiway delay with delays for arrival gates. Nevertheless, examining the data carefully does, we believe, give useful indications of effects of groundside congestion.

From the data for June 1995, we generated monthly averages of arrival and departure rates for each of 96, 15-minute intervals covering a day, for several airports. We considered airports widely known as busy ones, as well as one two not-so-busy terminals for comparison. We also generated monthly averages of taxi-out delays per flight, and of taxi-in delays per flight, for each of the 15-minute intervals in which arrivals or departures occurred at a rate of at least one per day. (Per-flight averages seem not likely to be meaningful when calculated for rare events.)

The results are shown in Figures 2-1 through 2-14. Peaks in mean taxi-out delays are quite reliably associated with peak departure rates, suggesting that the reported delays may be due to departure queues. The plots of taxi-in delays are quite often essentially flat. When peaks in taxi-in delays do occur, they are sometimes associated not with peaks in arrival rate, but with peaks in departure rate. This suggests that departure queues may, in fact, impact taxi-in operations. For example, taxi-in delay peaks are associated with departure peaks at The William B. Hartsfield Atlanta International Airport (ATL[1]) near 16:00 and at Chicago O'Hare International Airport (ORD) near 15:00, as well as with the morning taxi-in delay peak at Los Angeles International Airport (LAX). A peak in taxi-in delay may, however, be associated with a peak in arrival rate: the taxi-in peak at LAX near 21:00 gives an example of this.

Our PAMS departure data for Newark International Airport (EWR) apparently are corrupted. For completeness, we include a plot of departure delay that we derived from an Airline Service and Quality Performance (ASQP) source for June 1993, Figure 2-4.

Figure 2-1. Mean Departure Rate and Mean Delay in June 1995 for The William B. Hartsfield Atlanta International Airport

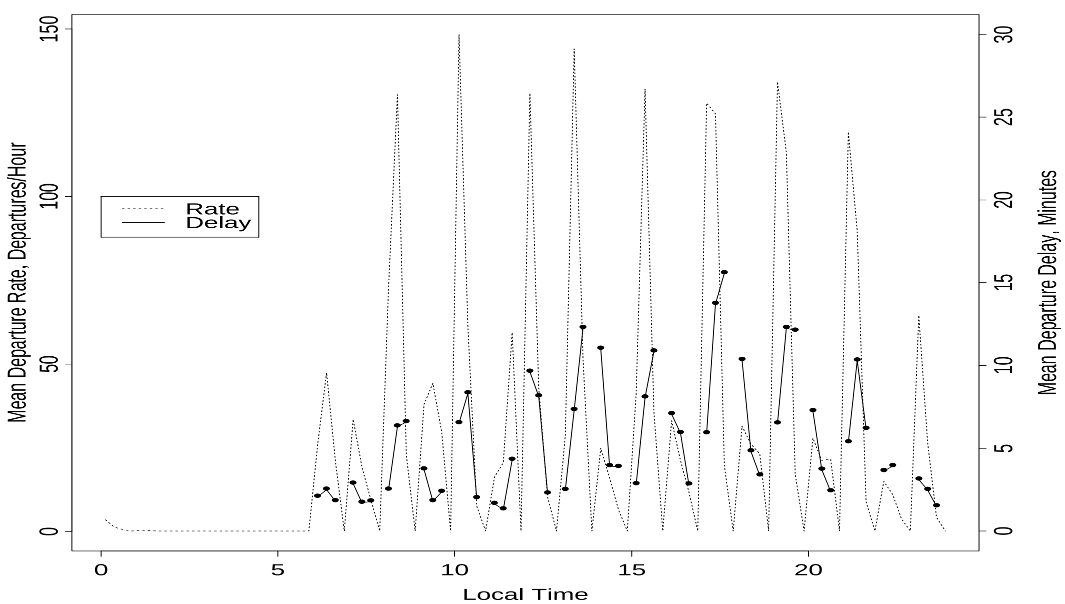

[1] Throughout the report, we refer to airports by their FAA identifiers. These are identified fully in the Appendix A, Glossary A Airport Identifiers, p. A-1.

Figure 2-2. Mean Departure Rate and Mean Delay in June 1995 for Baltimore-Washington International Airport

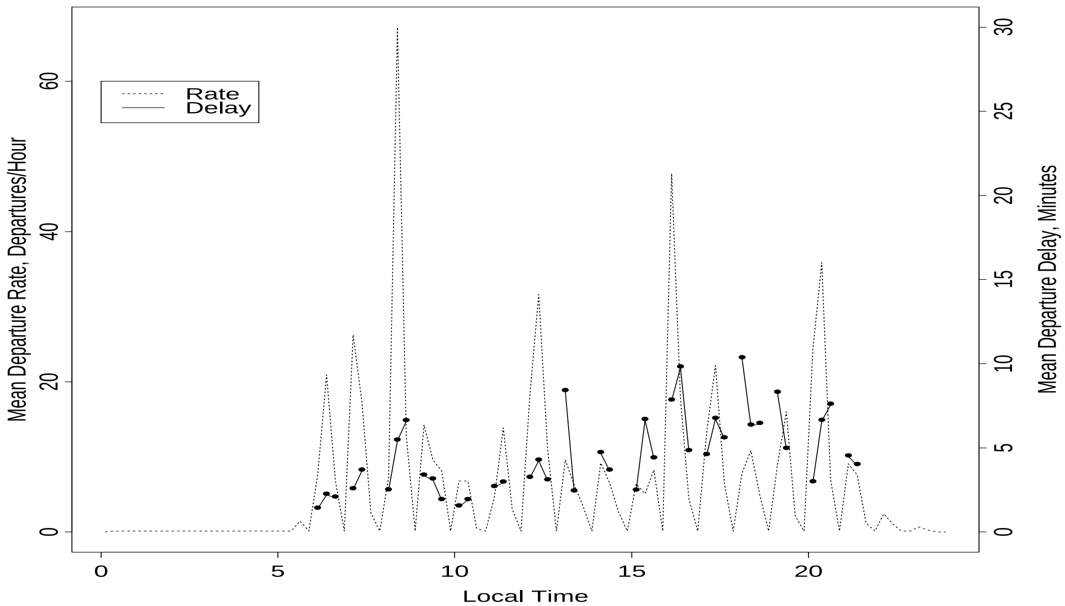

Figure 2-3. Mean Departure Rate and Mean Delay in June 1995 for Dallas-Fort Worth International Airport

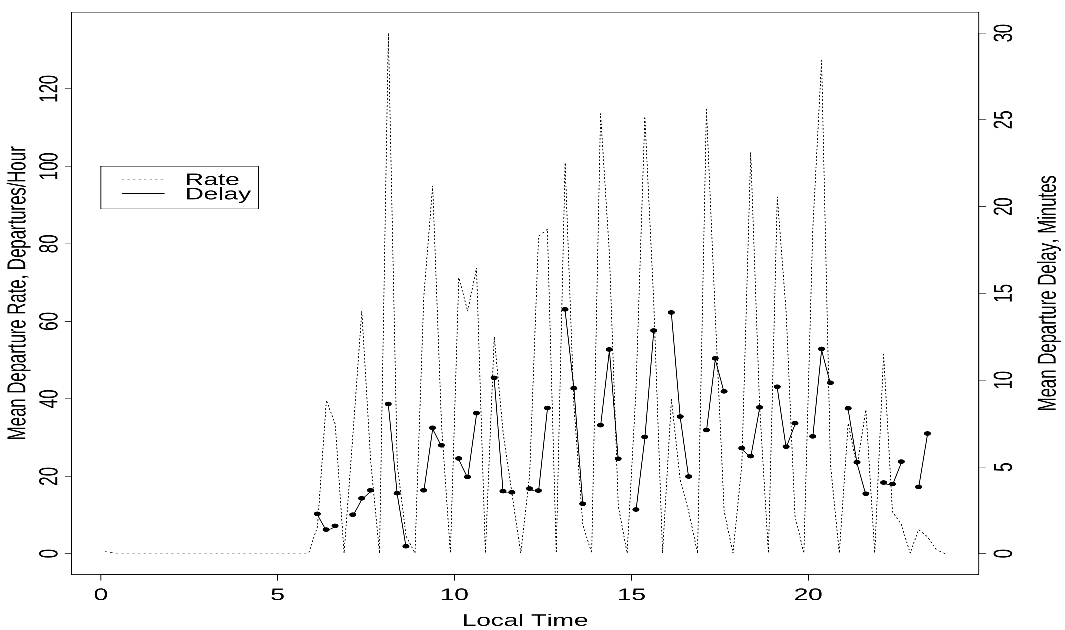

Figure 2-4. Mean Departure Rate and Mean Delay in June 1993 for Newark International Airport

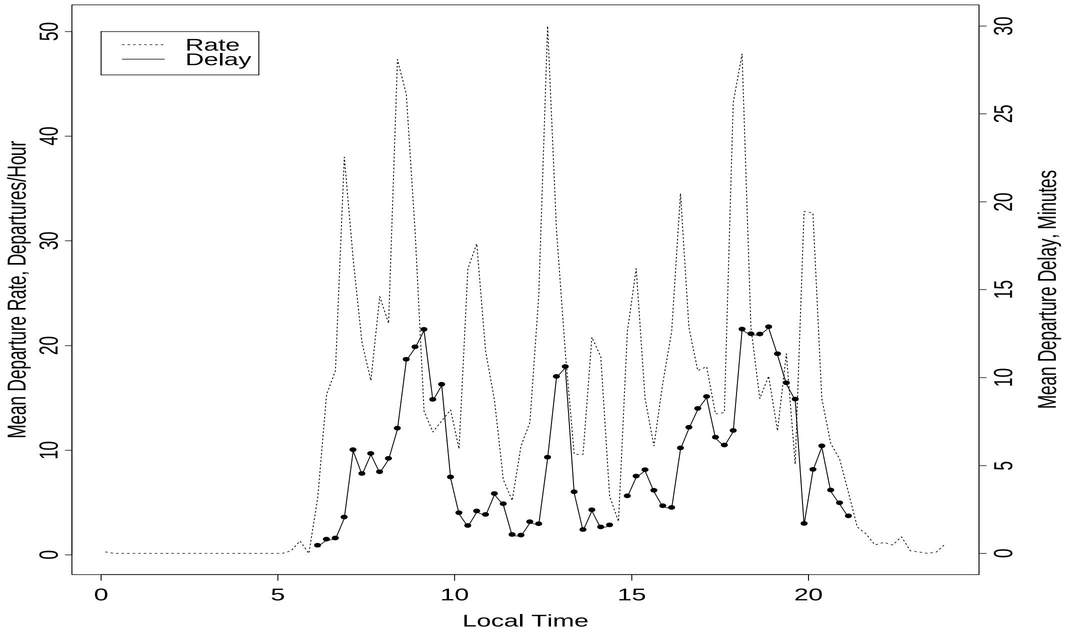

Figure 2-5. Mean Departure Rate and Mean Delay in June 1995 for Indianapolis International Airport

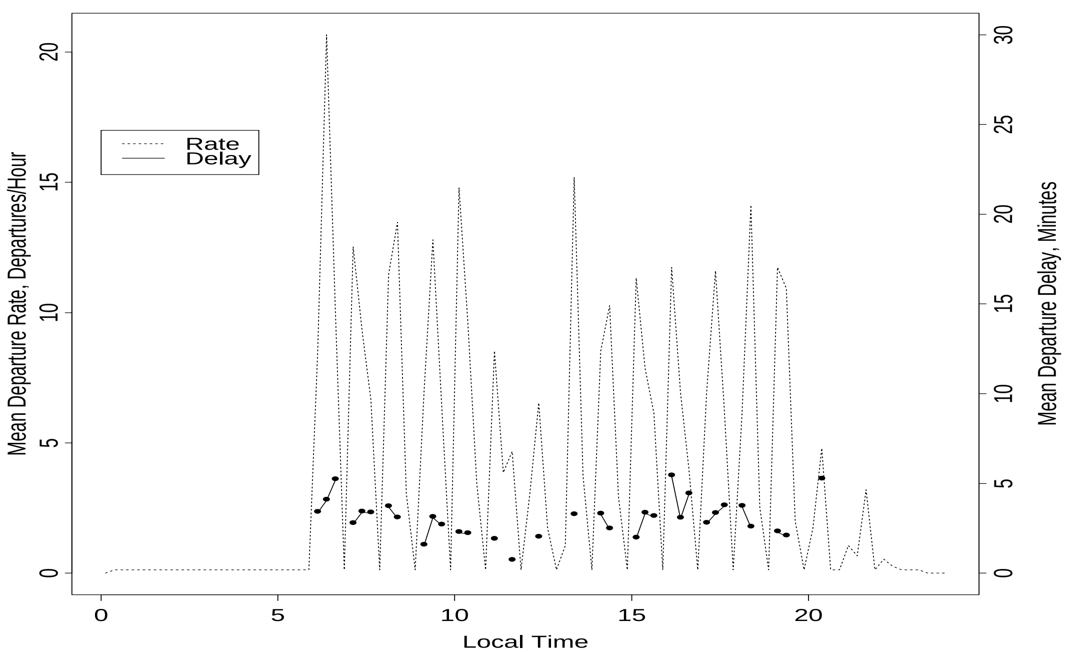

Figure 2-6. Mean Departure Rate and Mean Delay in June 1995 for Los Angeles International Airport

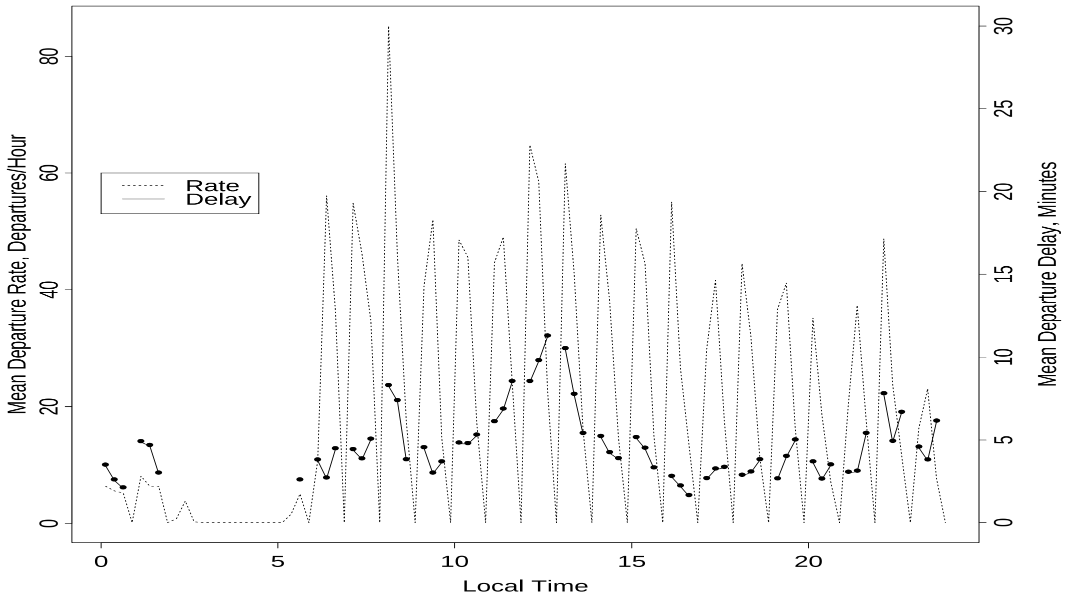

Figure 2-7. Mean Departure Rate and Mean Delay in June 1995 for Chicago O'Hare International Airport

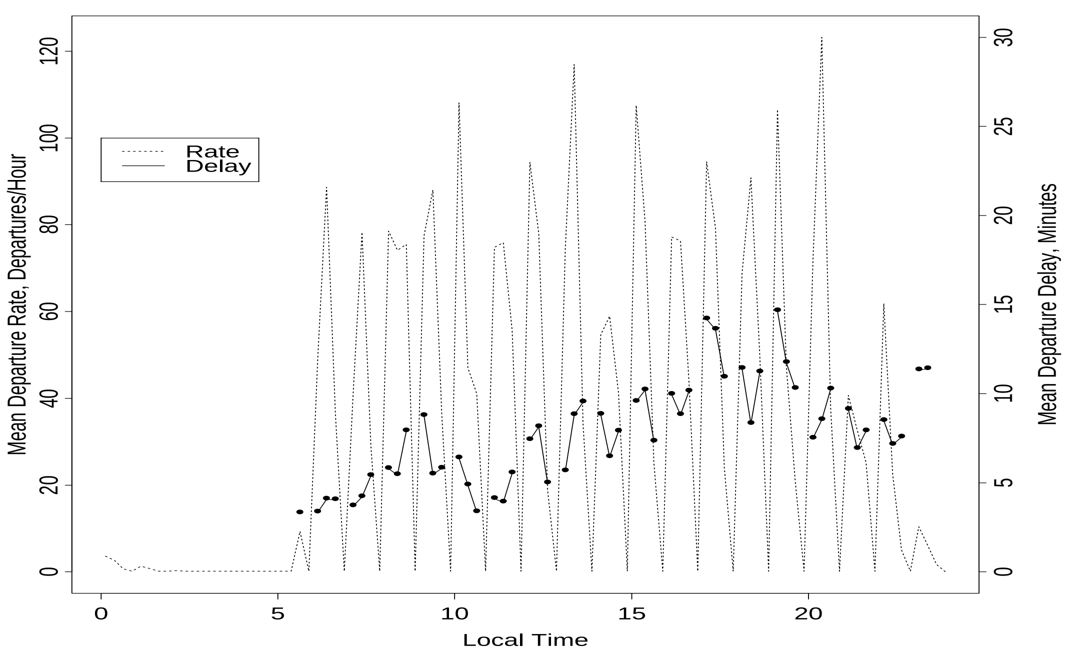

Figure 2-8. Mean Arrival Rate and Mean Delay in June 1995 for The William B. Hartsfield Atlanta International Airport

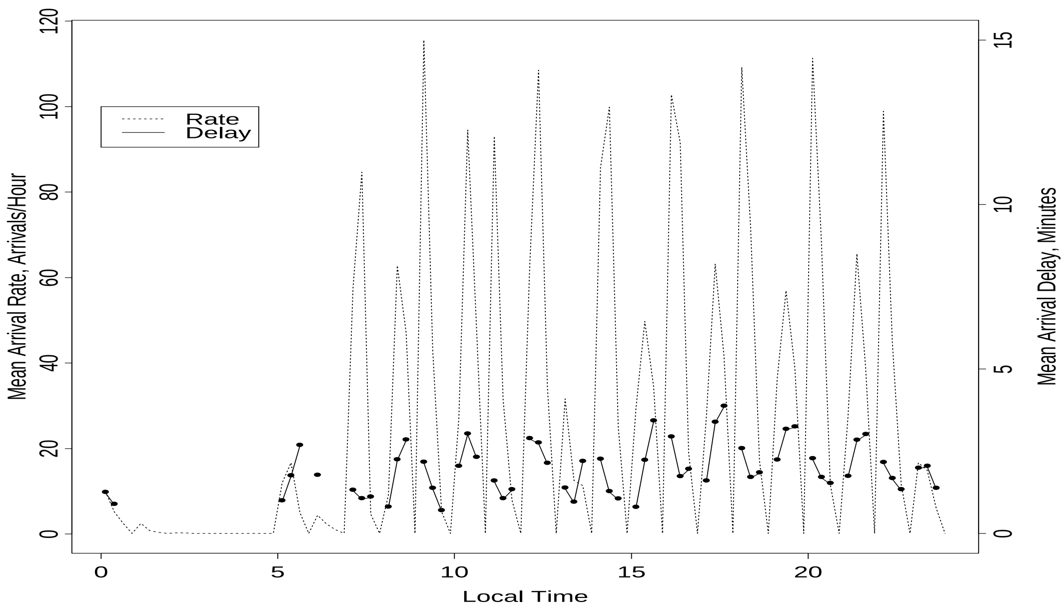

Figure 2-9. Mean Arrival Rate and Mean Delay in June 1995 for Baltimore-Washington International Airport

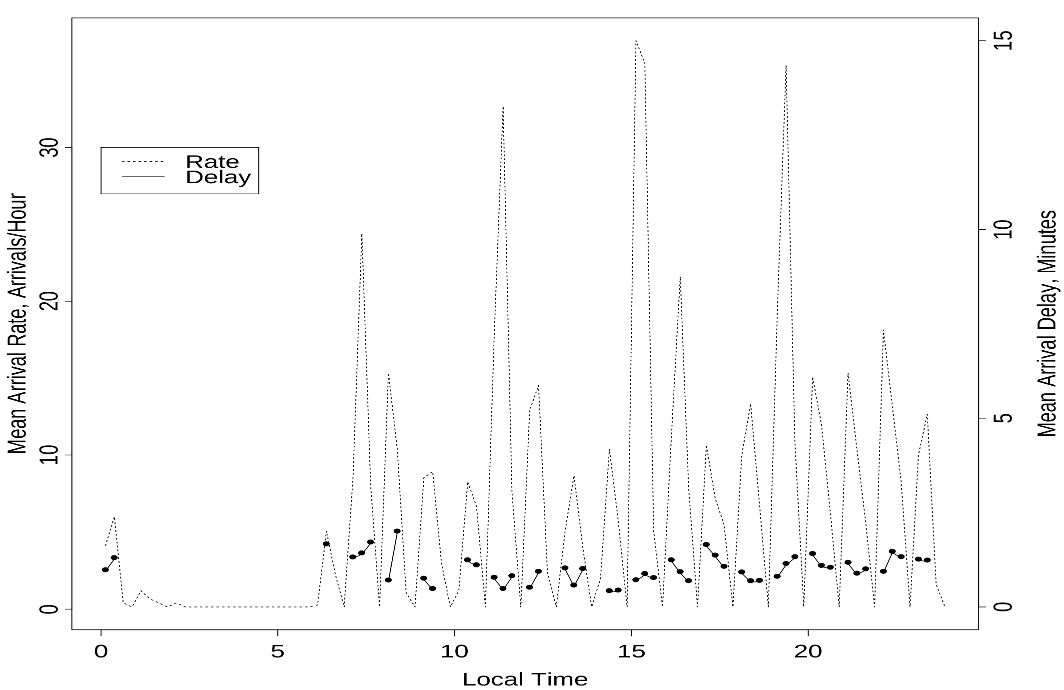

Figure 2-10. Mean Arrival Rate and Mean Delay in June 1995 for Dallas-Fort Worth International Airport

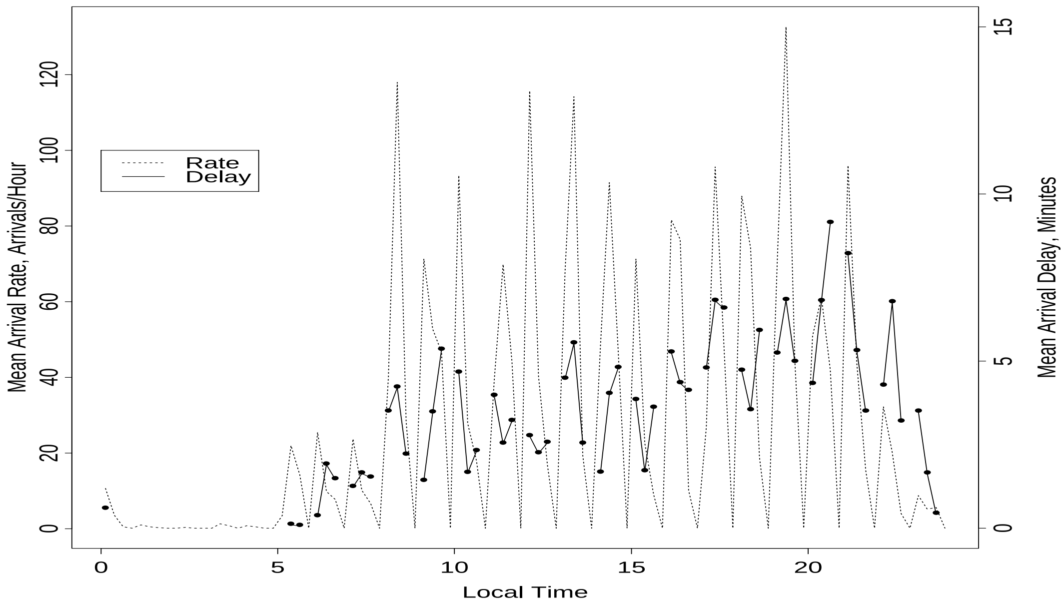

Figure 2-11. Mean Arrival Rate and Mean Delay in June 1993 for Newark International Airport

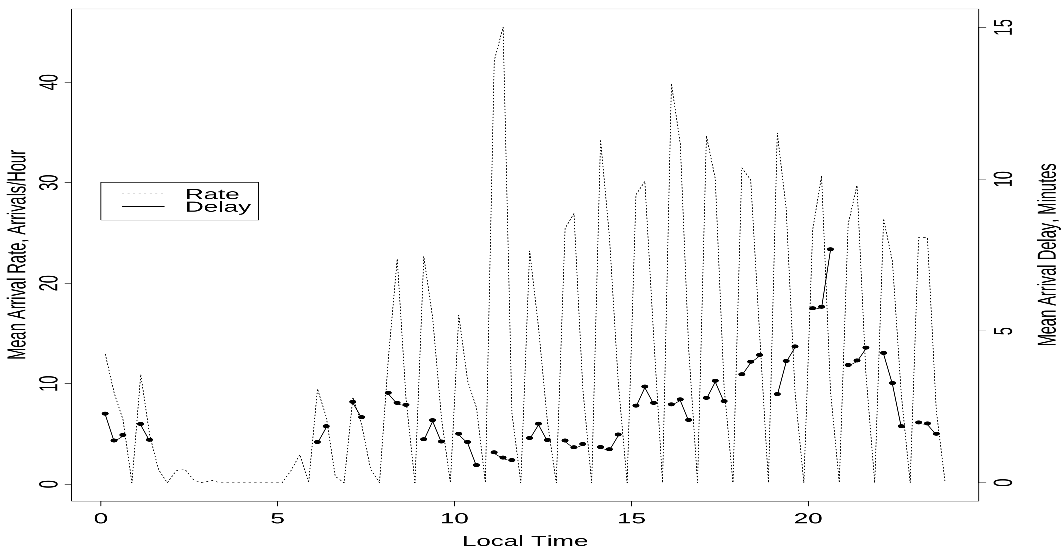

Figure 2-12. Mean Arrival Rate and Mean Delay in June 1995 for Indianapolis International Airport

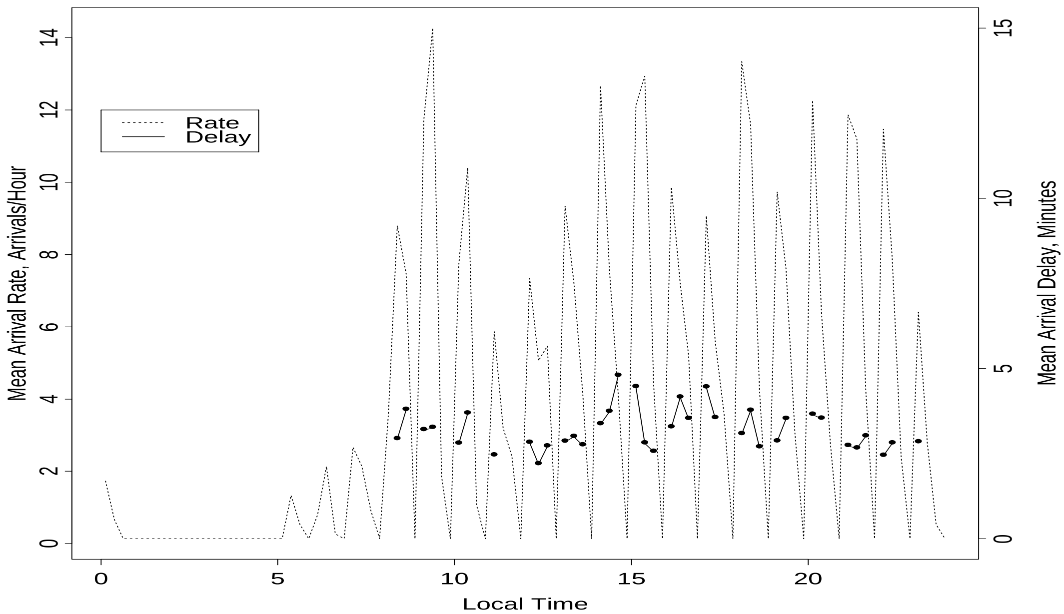

Figure 2-13. Mean Arrival Rate and Mean Delay in June 1995 for Los Angeles International Airport

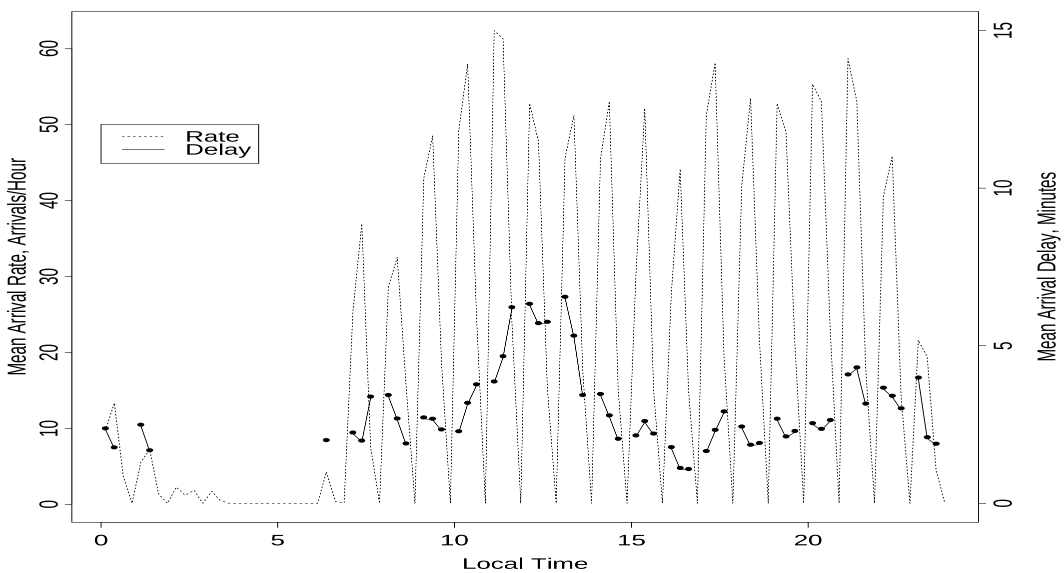

Figure 2-14. Mean Arrival Rate and Mean Delay in June 1995 for Chicago O'Hare International Airport

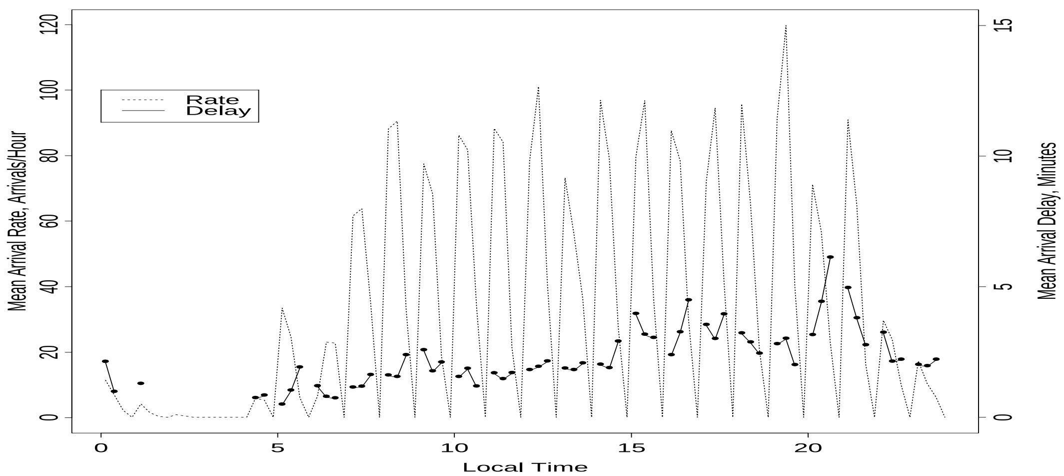

Taxi-out delays generally exceed taxi-in delays; overall, the seven airports we considered, the mean taxi-out delay per operation is slightly less than 6 minutes, and the mean taxi-in delay is just over two minutes. As one would expect, the busiest terminals generally show the greatest delays.

Table 2-1 summarizes mean taxi-in and taxi-out delays for June 1995, as reported in the FAA's Performance Monitoring Analysis Capability (PMAC) data system. These results generally are consistent with the delays shown in Figures 2-1 through 2-14.

Table 2-1. Mean Taxi-in and Taxi-out Delays

Airport	Average taxi-in delay per arrival (minutes)	Average taxi-out delay per departure (minutes)
ATL	2.2	9.3
BWI	0.8	3.6
DFW	4.3	5.1
EWR	1.7	12.0
IND	0.7	1.3
LAX	2.6	3.8
ORD	2.5	4.7
Mean	2.11	5.7

In interviews with both aircrew and controllers, we were told that most delays between gate-out and wheels-off were due to queues for departure service. We also were given anecdotes about queues for departure runways causing taxi-in delays when the departure queues became long enough to interfere with taxi-in runways.

Controllers and aircrew also mentioned taxi delays, when taxiways crossed active runways. Controllers also mentioned certain airports—ATL and DTW, for example—at which specific features caused taxi delays. For example, at DTW operations at one set of gates can interfere with taxiways.

Because of the confusion of departure runway queues with taxi-out delays and of gate delays with taxi-in delays, it is difficult, with the data available to us now, to compare costs of taxi-in and taxi-out delays with the costs of delays from other causes. Nevertheless, we wish to give some indication of the costs of taxi delays at major airports, to provide an indication of the opportunity for DSTs that improve ground movement.

We will do so this way: assuming that holds of taxi-in aircraft for lack of gate space are much less common than delays of taxi-out aircraft in queues for departure runways, we take the taxi-in delays as likely to be associated with actual taxiway congestion problems. Taking the mean taxi-in delays of less than 1 minute seen at the less-busy airports as a measure of what might be achieved with better surface movement management even at busy airports, we see potentially manageable delays of about 1 minute per operation at ATL, EWR, LAX, and ORD, and of about 3 minutes at DFW.

We take 1 minute per operation as a rough, and conservative, estimate of the taxi-in delay that may be eliminated by better surface movement management at busy airports. There seems to be no obvious reason to expect aircraft taxiing out to experience different effects of surface congestion than those seen by aircraft taxiing in, (although they do experience different effects on runway congestion) we also take this value as a rough estimate of the taxi-out delay that might be eliminated by improved surface movement management.

According to Department of Transportation data for 1994, the 12 busiest Contiguous United States (CONUS) terminals (in terms of operations per day) all had more than 750 operations per day. It seems reasonable to assume that at least these 12 airports experience surface congestion. Pricing a minute's taxi delay at $33.00 (see Table 3-5) leads to roughly $9 million per year of potentially avoidable taxi-delay costs at each such airport. Over just the top 12 airports, then, those costs amount to more than $100 million. We believe this is a conservative estimate of those costs; it appears ample to justify investing in the development of DSTs to help controllers make surface movements more efficient at busy terminals.

(As this report was in final editing, we received a report on observed benefits of the Surface Movement Advisor at ATL[2]. The report indicates taxi-time reductions averaging about 1 minute per operation, and annual savings at ATL of $16 million to $21 million.)

[2] Rada, Wilma, "Surface Movement Advisor (SMA) Benefit Analysis," MCA Research Corporation, Arlington, Virginia, October 1997.

INEFFICIENT ROUTES

This section reports results of our efforts to identify inefficiencies in the three-dimensional routing of aircraft in the NAS, that NASA DSTs could address. We primarily intended to identify opportunities rather than to evaluate them precisely; so we made mostly *qualitative* surveys. We did, however, develop some preliminary *quantitative* indicators of the relative importance of the opportunities that we found. We did not apply sophisticated survey methods, because there appeared to be no need for them: the inefficiencies were quite readily apparent.

To see how DSTs affecting routes between terminal areas might save operating costs, we examined two sources. Our principal source was ETMS data for two days, April 8, 1996, and November 27, 1996. A second very helpful source was a set of data from a major air carrier on the routes that they requested for certain Boeing 727 flights on August 1 through August 10, 1997.

In analyzing each source, we compared fuel burns on the cruise portions of the as-flown or as-requested routes, with fuel burns on optimal routes. We considered departure, cruise, and arrival phases of flight separately. The following sections give details of our methods and results.

Enhanced Traffic ManagementSystem (ETMS) Data

We examined ETMS data for Monday, April 8, 1996 and Wednesday, November 27, 1996. The April date was characterized by generally good weather over most of the CONUS, with significant periods of Instrument Meteorological Conditions (IMC) at certain terminals. A cold front running from Seattle to just north of San Francisco brought showers and reduced visibility to the northwestern corner of the CONUS. Another cold front pushing southward across the Appalachians produced showers and snow flurries in the northeast, notably at Boston, and into the mid-Atlantic states. A stationary front caused showers and some reduced visibility over the Florida peninsula. While there were some locally significant delays like at Boston, air traffic over the bulk of the CONUS seems not to have been significantly disrupted and ETMS arrivals and departures at most of the 64 LMINET airports generally matched the OAG.

The November date was, of course, the day before Thanksgiving, typically one of the busiest air travel days of the year, if not the busiest.

CRUISE ROUTES

This section reports results of our comparison of ETMS as-flown routes, with optimal routes. As an illustration, Figure 2-15 compares some ETMS routes flown from LAX to BOS on April 8, 1996 with an optimal trajectory.

Figure 2-15. Example Cruise Routes

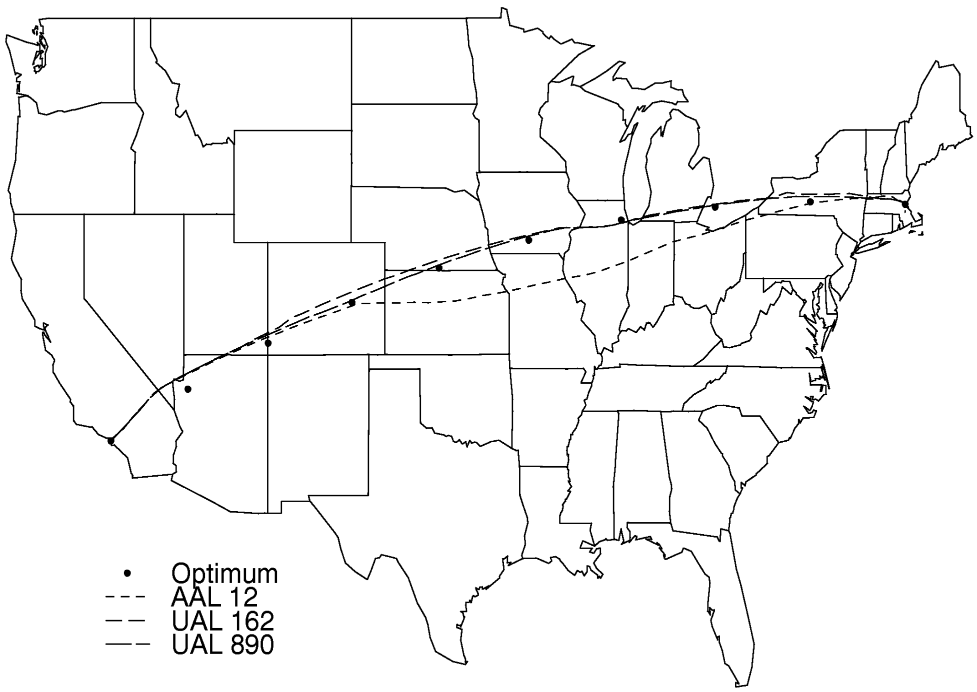

The routes of three flights are shown: AAL12, UAL162, and UAL890. The dots indicate the wind route (optimal route) for April 8, 1996. Only one flight, AAL12, departed significantly from the optimal route.

To examine the efficiency of the cruise portions of a selection of ETMS records, we identified top-of-climb (TOC) and top-of-descent (TOD) points visually. Automatic recognition of these points is probably possible, but, as Figure 2-16 shows, altitude holds during climb and descent, and step climbs en route, complicate the task.

Figure 2-16. Vertical Profile of Flight from JFK to LAX

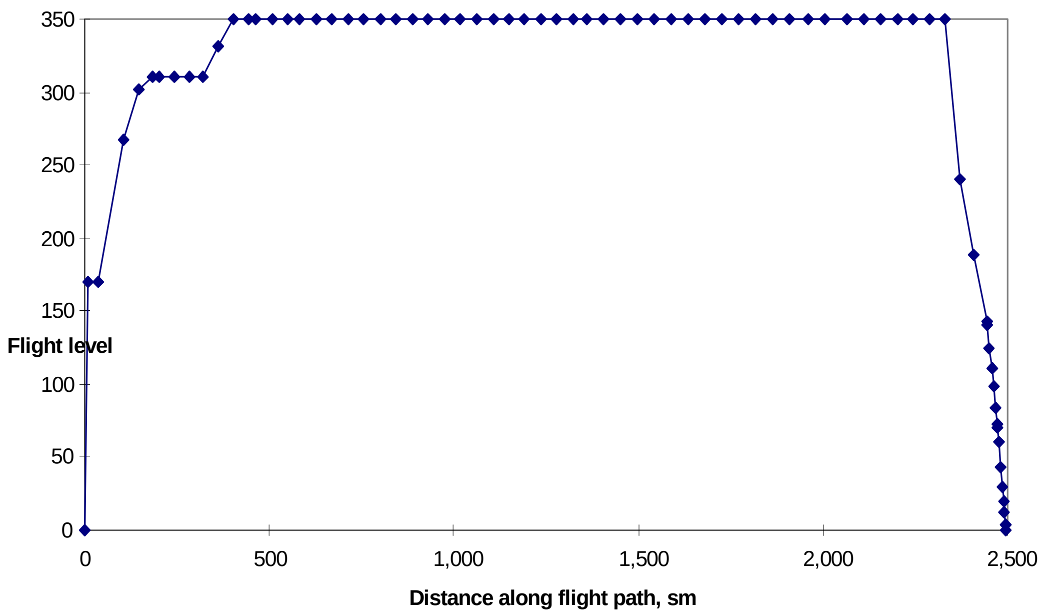

We determined the fuel burn on the as-flown route, between Top of Climb (TOC) and Top of Descent (TOD), for winds aloft during the date of the flight. We then found the fuel burn on the optimal trajectory between TOC and TOD. Table 2-2 shows the results.

Table 2-2. Savings from Flying Optimal Cruise Routes Vice ETMS Cruise Routes

Origin	Destination	As flown fuel, LB	Time (hour)	Optimal fuel, LB	Time (hour)	Fuel saved (percent)	Time saved (min)
BOS	LAX	34,047	4.71	33,813	4.68	0.69	1.8
EWR	LAX	29,524	4.26	29,418	4.25	0.36	0.6
SEA	LAX	14,292	2.06	14,192	2.05	0.70	0.6
ORD	LAX	18,745	2.73	18,596	2.71	0.79	1.2
YVR	LAX	11,828	1.93	11,785	1.92	0.36	0.6
DFW	LAX	11,534	1.62	11,502	1.62	0.28	0
IAD	LAX	25,792	3.64	25,740	3.63	0.20	0.4
MIA	LGA	11,001	1.79	10,935	1.78	0.60	0.6
BOS	LAX	34,047	4.71	33,813	4.68	0.69	1.8
EWR	LAX	29,524	4.26	29,418	4.25	0.36	0.6
SEA	LAX	14,292	2.06	14,192	2.05	0.70	0.6
ORD	LAX	18,745	2.73	18,596	2.71	0.79	1.2
YVR	LAX	11,828	1.93	11,785	1.92	0.36	0.6
DFW	LAX	11,534	1.62	11,502	1.62	0.28	0

Table 2-2. Savings from Flying Optimal Cruise Routes Vice ETMS Cruise Routes (Continued)

Origin	Destination	As flown fuel, LB	Time (hour)	Optimal fuel, LB	Time (hour)	Fuel saved (percent)	Time saved (min)
IAD	LAX	25,792	3.64	25,740	3.63	0.20	0.4
MIA	LGA	11,001	1.79	10,935	1.78	0.60	0.6
LAX	EWR	27,707	3.96	27,640	3.95	0.24	0.6
SEA	EWR	28,307	3.60	28,284	3.60	0.08	0
ATL	EWR	3295	0.52	3,281	0.52	0.42	0
DCA	DFW	11,708	1.77	11,673	1.77	0.30	0
LAX	DFW	13,848	1.92	13,816	1.92	0.23	0
DEN	DFW	3,588	0.60	3,583	0.60	0.14	0.1
BNA	DFW	4,648	0.72	4,648	0.72	0.00	0
BOS	DFW	17,134	2.56	16,570	2.48	3.29	4.8

In all but one case, fuel savings were less than 1 percent. That level of difference might be accounted for by differences between our winds-aloft data and the wind forecasts available to the airline operations centers.

The exceptional case turns out to be a revealing one. Figure 2-17 shows the route of that flight, with some streamlines for winds aloft at its cruising altitude of 39,000 feet.

Figure 2-17. BOS-DFW Flight Path and Winds Aloft

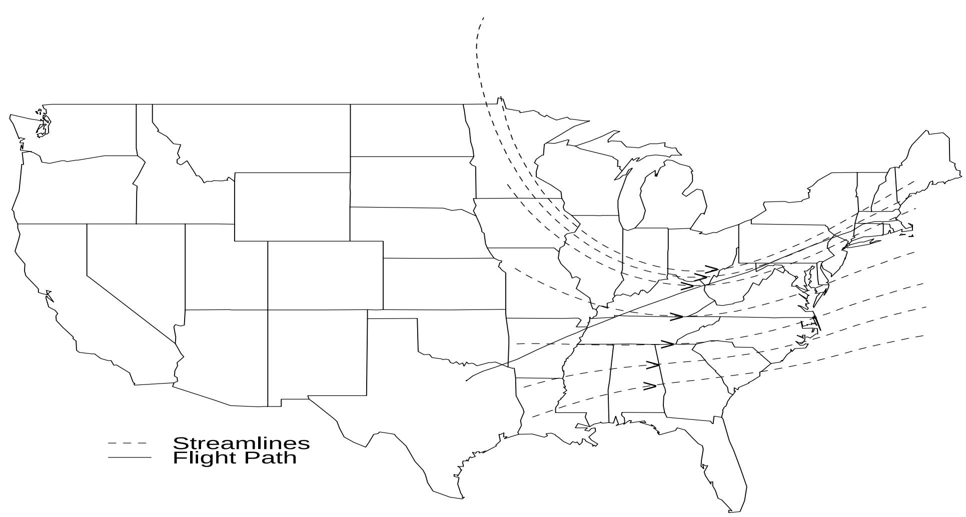

The streamlines show that much of the flight path passed along the edge of the front that was advancing into the North Atlantic states. One would expect that to be a region of rapidly varying winds, and it was. Figure 2-18 shows headwinds at the flight's cruising altitude, at points on a line perpendicular to the ETMS path located about one-third of the way from BOS to DFW. The points ranged from 100 miles southeast of the ETMS path through 100 miles northwest of the path. For the conditions experienced by this flight, displacing its path by only a few tens of statute miles makes a substantial difference in headwinds, and thus in fuel burn.

In view of the turbulence to be expected near a front, it is by no means clear that the minimum-fuel cruise route would have been an acceptable one.

Overall, our survey of ETMS trajectories did not discover widespread examples of inefficient cruise paths. In the one exceptional case that we found, it was not clear that the optimal cruise path would have been practical.

Figure 2-18. Head Winds for BOS-DFW Flight Versus Distance from ETMS Path

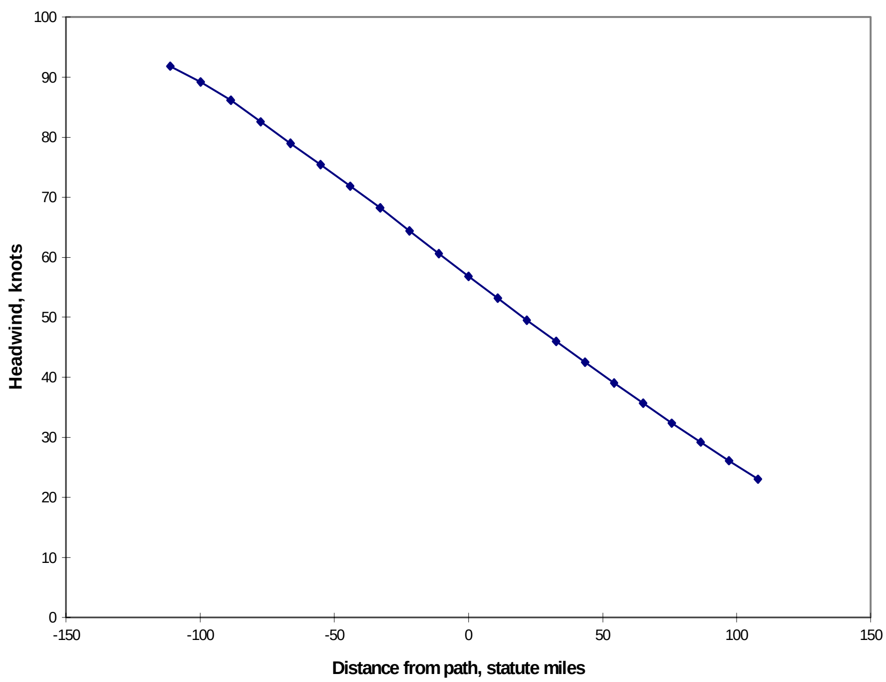

CRUISE ALTITUDES

Fuel burns certainly vary with altitude, and airlines do complain about not getting desired cruise altitudes. In this subsection, we report on work to quantify the benefits to be expected from this aspect of cruise. Figure 2-19 compares the as-flown altitude profiles for three Boeing 757 flights from LAX to BOS, with an optimal profile. The flights took place on April 8, 1996, and the optimal profile is for the winds of that date, for a Boeing 757-200.

Figure 2-19. Comparison of Altitude Profiles

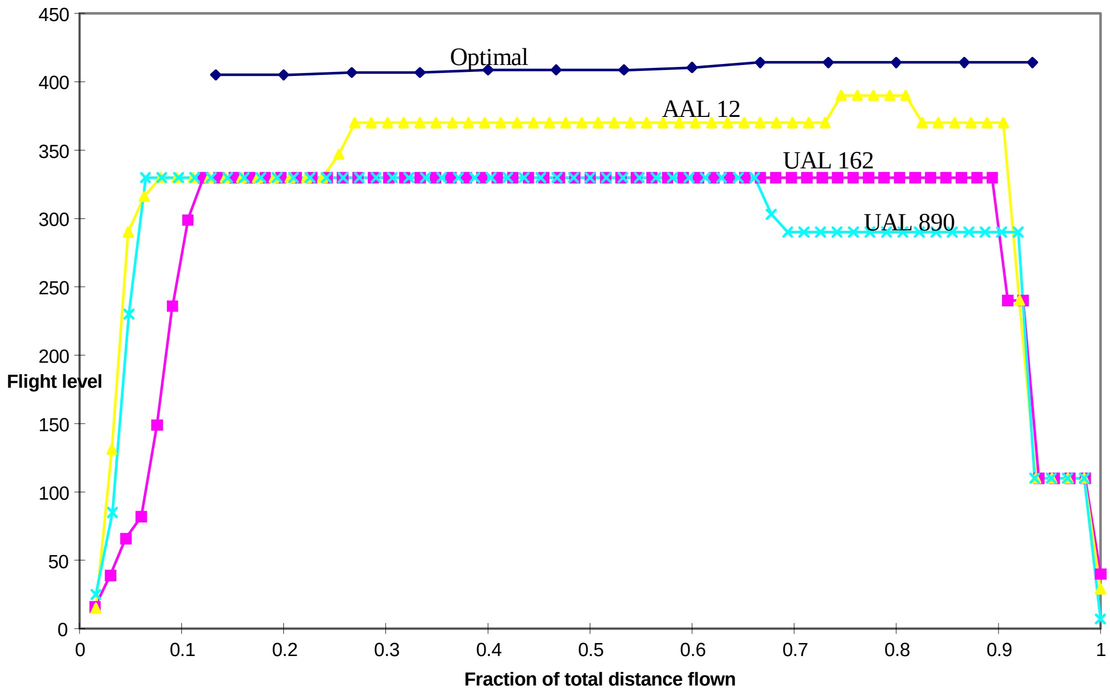

The optimal profile reached altitudes near the maximum for the aircraft. None of the flights cruised as high as our optimal profile, although AAL12 flew close to it for a relatively short part of the trip.

The differences between our optimal cruise altitude and those flown by the sample flights are not necessarily due to ATM constraints. While aircraft type and winds aloft are primary determinants of best-cruise altitude, these data alone do not specify the desirable cruise altitude for a given flight. The optimal cruise altitude for a given aircraft varies with takeoff weight. Dispatchers generally seek minimum-fuel trajectories consistent with schedule integrity. This produces considerable situational variation in altitude and Mach number profiles.

Goals other than fuel economy often dictate the choice of cruise altitude. Avoiding turbulence often determines the altitude flown. An experienced controller at the FAA's Command Center told us, however, that 60 percent to 80 percent of flights are *not* affected by turbulence.

The many factors affecting cruise altitude seem to make direct comparisons between ETMS cruise altitudes and optimal ones fruitless. Nevertheless, it seems worthwhile to explore the scale of the inefficiencies that result from operating flights away from their optimal altitudes. The variation of fuel burn with altitude for a Boeing 757 cruising for 1,000 nautical miles in calm winds is shown in Figure 2-20. To obtain the results in Figure 2-20, we varied Mach number with altitude to keep cruise time constant at 2.25 hours, and the fuel cost of climb from 29,000 feet to cruise altitude is included in the fuel burns shown.

Figure 2-20. Variation of Fuel Burn with Cruise Altitude

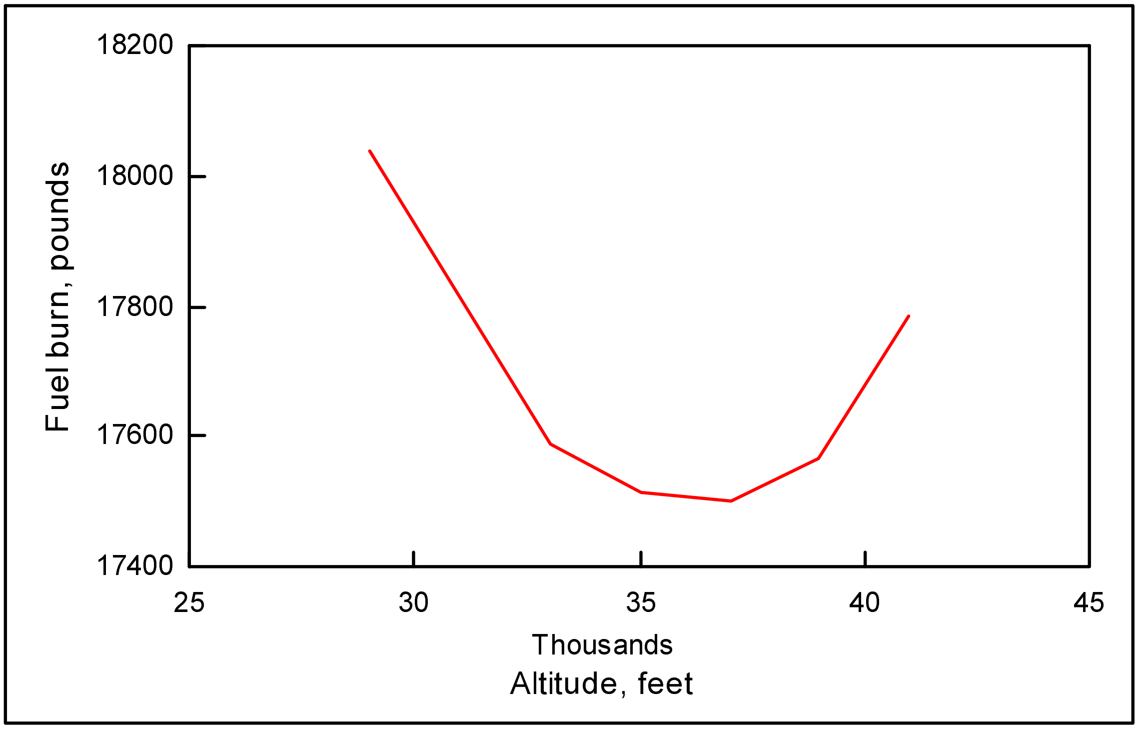

The optimal cruise altitude for the case illustrated in Figure 2-20 is 37,000 feet. That altitude is not available to westbound traffic in the FAA's IFR altitudes. If the flight had been westbound, the closest available altitudes would have been 35,000 feet and 39,000 feet. Operating at either of those altitudes would have cost relatively little in added fuel: 15 pounds at 35,000 feet and 69 pounds at 39,000 feet.

If the flight were eastbound, and were for some reason denied the optimal altitude of 37,000 feet, the closest IFR altitudes would be 33,000 feet and 41,000 feet. Operating at those altitudes instead of at the optimal altitude would cost 89 pounds of fuel and 286 pounds of fuel, respectively. All the increments in cruise fuel are less than 1 percent except for operating at 41,000 feet, where the increased fuel burn amounts to 1.6 percent.

A Major Airline's Requested Routes

United Airlines gave us data on their planned routes between several CONUS cities, for August 1 through 10, 1997. Comparing fuel burns on these planned routes with fuel burns on optimal routes shows that, in some cases, operating on optimal routes can reduce block fuel by roughly 1 to 2 percent. Such savings are, however not common. The data also show that constraining flights to follow Standard Instrument Departures (SID) and Standard Terminal Arrival Routes (STARs) significantly reduces opportunities to reduce fuel burns by flying optimal trajectories.

An example of a case in which there were significant savings. For Boeing 727 flights from DEN to ORD, on four occasions, the carrier planned the Plains One SID from DEN, with the Hayes Center transition (HCT). The requested route continued over OBH and FOD, and joined the Janesville Four standard arrival route (STAR) to ORD at DBQ. As one can see, this route goes well to the north of the great circle. On August 3, 1997, the wind route from DEN to ORD also lay north of the great circle, but not so far north as the carrier's requested route.

Figure 2-21. Routes from DEN to ORD

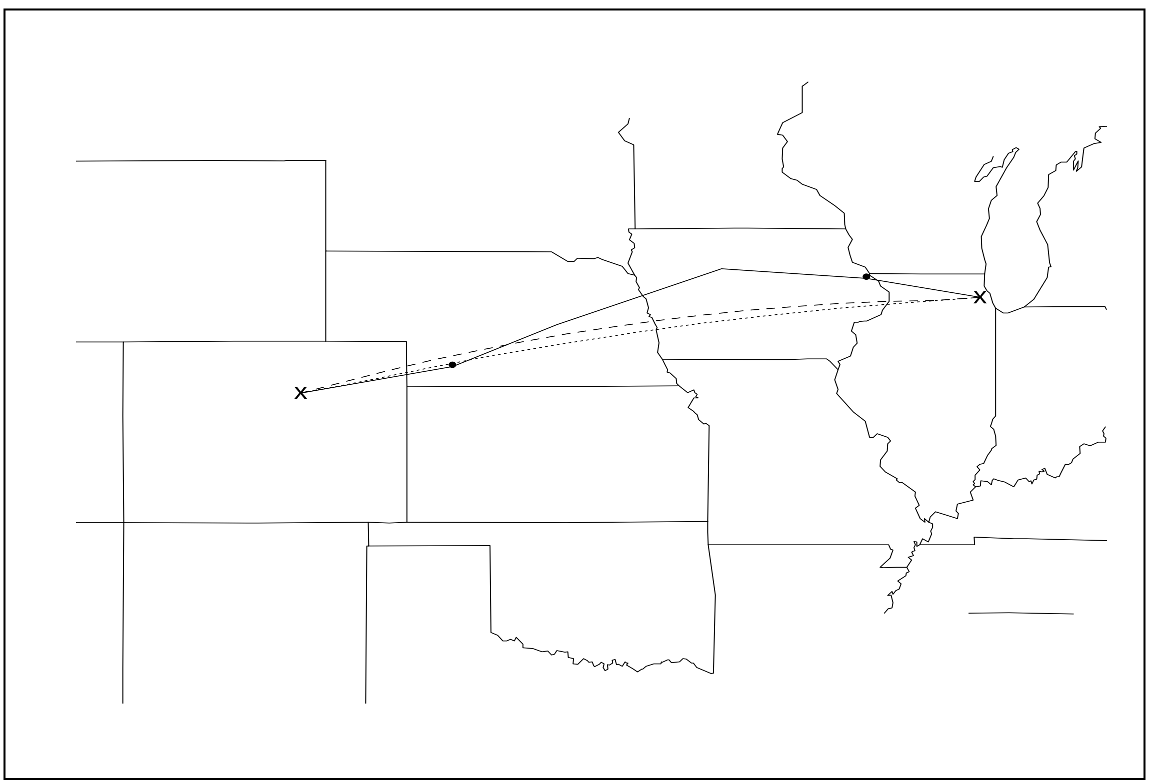

Note: Dotted curve is great circle, dashed curve is wind route for 4/3/96, and solid curve is carrier's requested route. Solid circles indicate end of SID and beginning of STAR.

Operating on the wind route would have saved 312 pounds of fuel, about 1.6 percent of block fuel. Constraining the flight to follow the SID through HCT, and to join the STAR at DBQ, reduced the optimal-route savings to 166 pounds, about nine-tenths of 1 percent. Average block fuel savings over the 4 days were 320 pounds (1.65 percent) for unconstrained wind routes, and 167 pounds (0.85 percent) for wind routes that could be adjusted only between HCT and DBQ.

TERMINAL-AREA AIRSIDE CONGESTION

We considered both departures from, and arrivals to, several terminal areas. As detailed in the following subsections, we found many examples of inefficient arrivals at busy terminals. We found fewer examples of inefficient departures.

Departures

We examined both altitude profiles and routes for indications of inefficiencies for departures from congested airports and others. As one would expect, departures from the New York area during busy periods exhibit both inefficient routes and altitude profiles. Figure 2-16 shows a flight departing JFK at 8:30 a.m. held at 17,000 feet for several minutes. Such holds are common at NYC during busy periods; Figures 2-22 and 2-23 show, respectively, examples of a departure from EWR at 6:04 p.m., and a 8:30 a.m. departure from LGA, respectively.

Figure 2-22. Altitude Profile of an EWR Departure

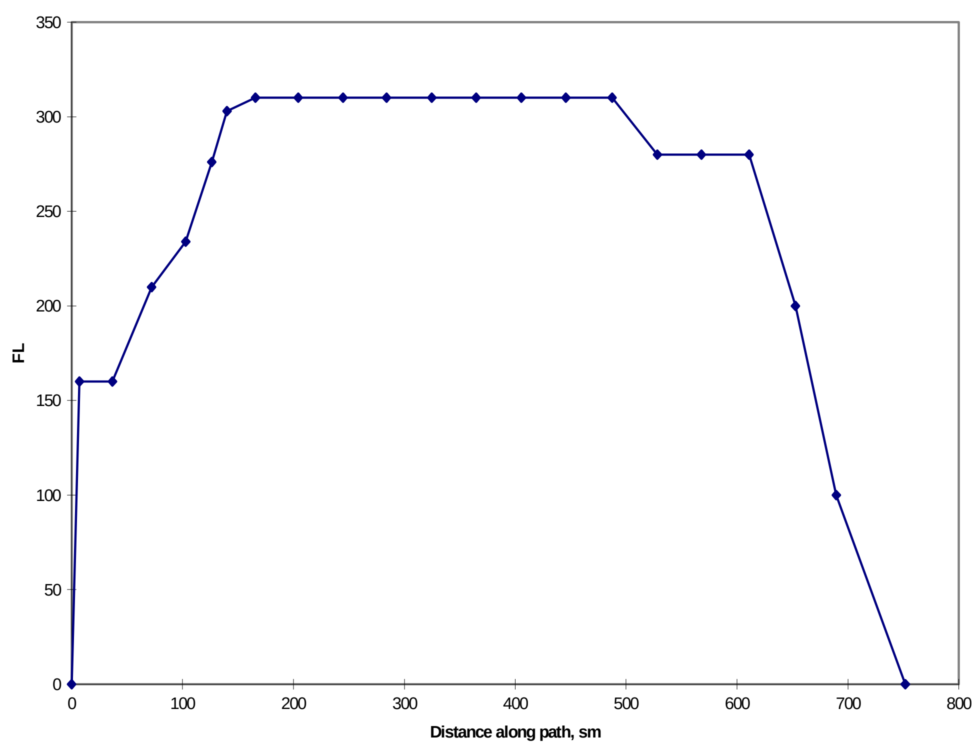

Figure 2-23. Altitude Path of LGA Departure

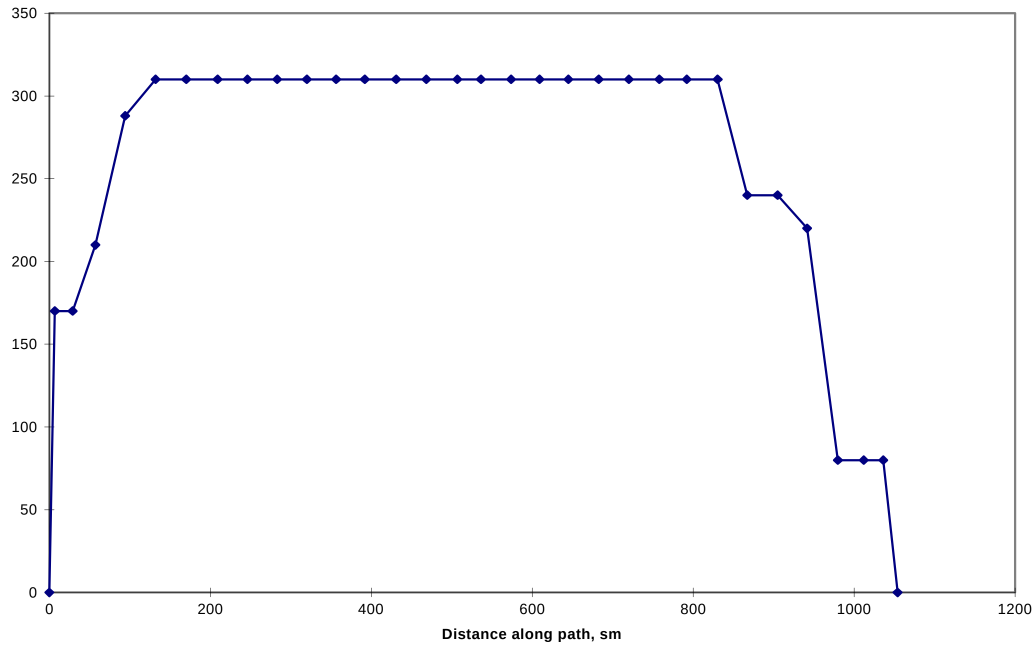

The paths taken by these three sample flights, which were bound for Los Angeles, Chicago, and Miami also obviously were constrained by traffic. Figure 2-24 illustrates this.

Figure 2-24. Departure Paths from NYC

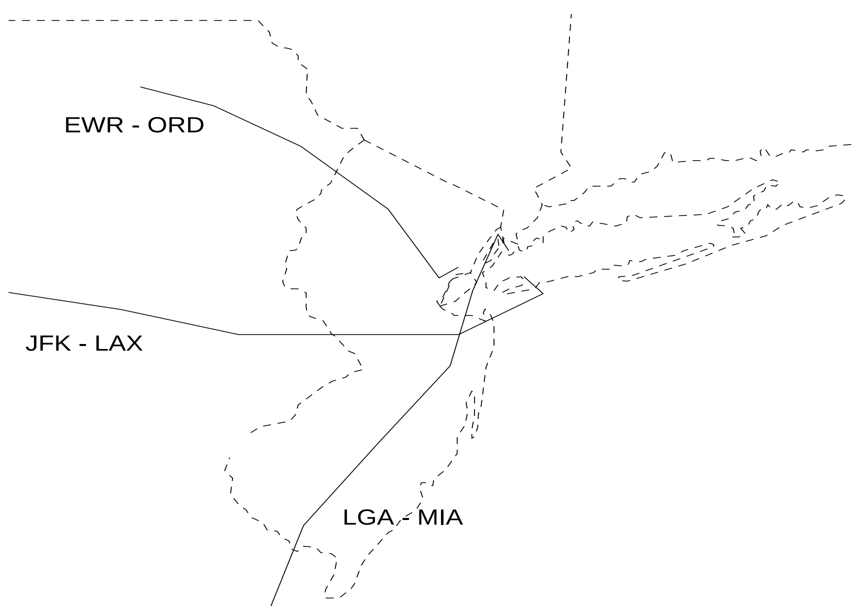

For a quantitative description of the efficiency of the departure process at NYC airports, we compared times to reach 30,000 feet for flights of more than 400 nm originating at IND and at the NYC terminals. The overall mean time for NYC departures was 16.6 minutes. Mean times to 30,000 feet for IND departures on flights of more than 400 nm was 13.1 minutes, which suggests that the less-congested conditions at IND enable generally more efficient departures. A standard t-test gives 95 percent confidence that the mean time for departures from NYC airports exceeds the mean time for departures from IND by 2.2 minutes.

The means for EWR and JFK departures do not vary significantly from the overall mean; the standard t-test gives 95 percent confidence that the mean time for LGA is 0.6 minutes less than the time for EWR and JFK. Table 2-3 shows times-to-climb at several terminals. The 3-minute difference between busy and not busy airports seems persistent.

Table 2-3. Mean times-to-climb

Terminal	Mean time to FL 300, min.
JFK & EWR	17.1
LGA	15.8
NYC	16.6
IND	13.1
ATL	16.3
ORD	15.9
MDW	15.2
CHI	15.8
OMA	13.9
All busy	16.2
All not-busy	13.5

Turbojet transport aircraft burn, roughly, 100 pounds of fuel per minute at low altitudes. Thus, the roughly 3-minute added time-to-climb for departures from busy terminals adds about 300 pounds to typical fuel burns.

Arrivals

Figure 2-23 illustrates an inefficient descent profile for an arrival at ORD. Such profiles are very common for arrivals at busy terminals. One quantitative measure of the efficiency of the arrival process is the time arriving flights spend at altitudes below those of efficient cruise-descent paths.

It seems likely that the initial descent shown in Figure 2-23, from cruise altitude to an altitude below FL 250, while still more than 200 miles from the destination, was forced by ATM and not desired by the crew. Such early descents are very common features of operations into busy terminals. Assessing this source of inefficiency quantitatively and fairly would, we believe, require some means of avoiding counting early descents that *were* requested by the crew, for example, in search of a smoother ride.

Interviews with controllers and aircrew suggest that 60 to 80 percent of flights operate without turbulence impact and that most turbulence is associated with obvious weather features such as fronts. Nevertheless, a significant fraction of flights clearly are impacted by turbulence; not all turbulent regions can be identified from gross features of synoptic weather.

In view of this, we decided that an adequate effort to identify turbulence effects would not be a reasonable and balanced use of this study's resources. Another feature of the arrival in Figure 2-23, the significant time spent below FL 100, is also quite commonly found for arrivals at busy terminals. This important source of inefficiency can be analyzed with available data.

Examining ETMS tapes, we found what appear to be significant variations in the amounts of time arriving flights spend at or below 11,000 feet. Table 2-4 displays results for April 8, 1996. (Flights of less than 400 nm are excluded, because in many cases, for example, flights from ORD to IND, the entire flight is at low altitudes.)

Table 2-4. Statistics of the Time Arriving Flights of More Than 400 nm Spend at or Below 11,000 Feet

Airport	<timelow>	s.d. timelow
EWR	19.2	3.2
JFK	8.2	10.4
LGA	10.8	11.7
NYC	14.3	9.7
DFW	20.0	7.1
LAX	17.0	1.6
IND	11.4	5.8
CHI	16.6	7.5
ORD	16.1	5.8
MDW	19.4	14.3

A reasonable approach would require some 10 or 12 minutes of flight at or below 11,000 feet. This value represents the times actually spent at a non-congested airport like IND. Some busy terminals also achieve those times, but others do not. The mean time below 11,000 feet for EWR, DFW, LAX, ORD, and MOW is more than 7 minutes greater than 11 minutes.

The differences among the three NYC airports may be explained by Visual Meteorological Conditions traffic patterns. As shown by Figure 2-25, arrivals to JFK (seen as traces beginning near the bottom center and coalescing into two streams ending to the right of the center of the frame) are given rather more straightforward approaches than are arrivals to the other airports. Arrivals to EWR (seen as traces beginning near the figure's lower left-hand corner, arcing up to the upper left-hand corner and then turning downward to the sharply-delineated final approach tracks above and to the left of the figure's center) appear to have the most circuitous paths of those for all the NYC airports, and this is consistent with the longer times of flight at lower altitudes seen for EWR.

Figure 2-25. Traffic Flows in NYC Area

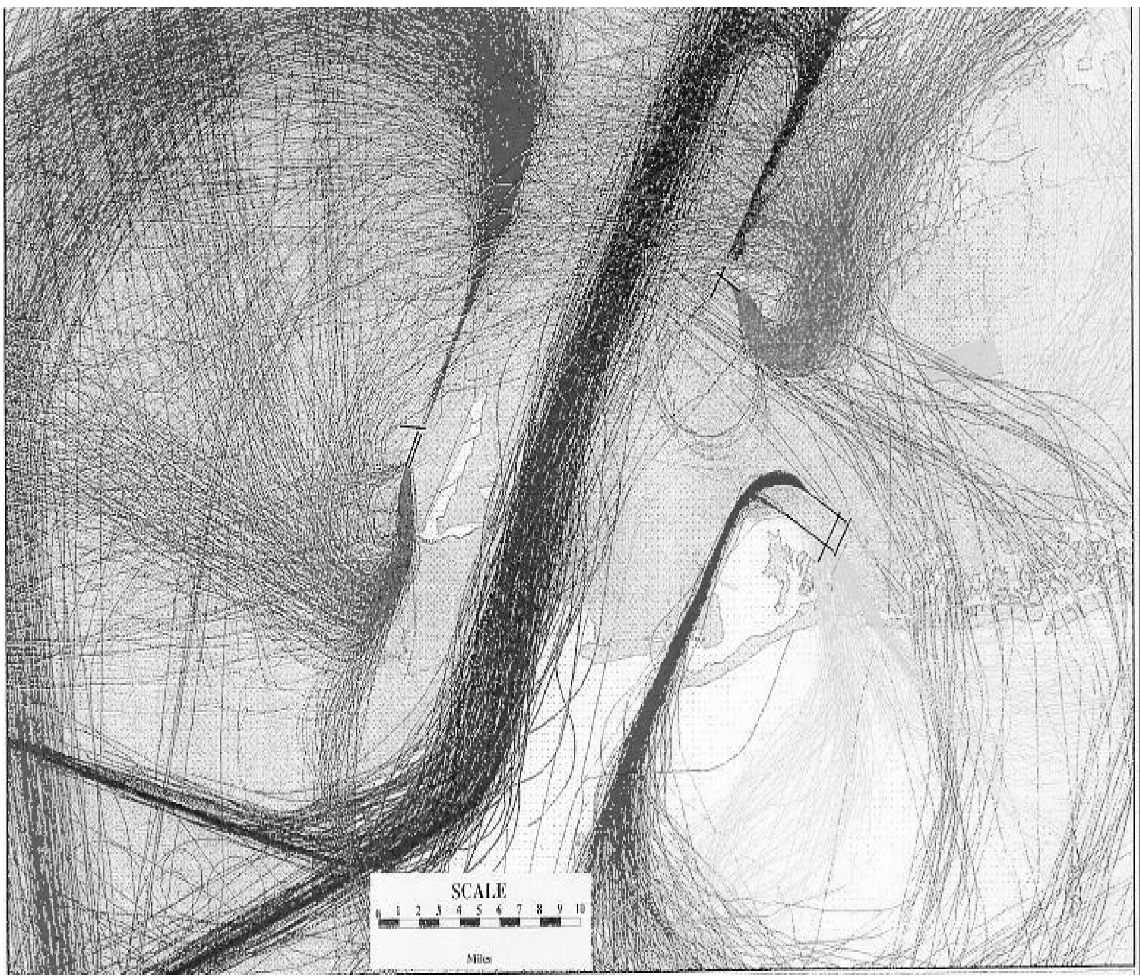

Large turbojet transports (B737 and MD80) burn roughly 500 pounds of fuel flying at low levels for 70 minutes. The Boeing 757 and the Boeing 727 each burn roughly 800 pounds and the MD11 burns roughly 1,400 pounds during such a flight. Our interviews with aircrew members showed that they were thoroughly aware of the losses on this scale caused by low-level operations during arrivals. Evidently, inefficient arrivals at busy terminals afford opportunities for NASA DSTs to effect significant savings.

A Suggestion From an FAA Controller

During a review of this project at the FAA Command Center, a controller made a suggestion for a DST that seems worth considering. He observed that not all ARTCCs and TRACONS are equally efficient in bringing traffic into holding patterns, and out of them. A DST that would advise controllers on efficient management of holding patterns would, he believed, have considerable potential.

Summary

Our review of opportunities for NASA DSTs to relieve inefficiencies in present NAS operations addressed departure, cruise, and arrival phases of flight, and airport groundside congestion. It identified inefficiencies in departures costing about 300 pounds of fuel per turbojet flight at busy terminals and inefficiencies in arrivals costing roughly 800 pounds of fuel per turbojet arrival at such terminals. Inefficient routes between terminals (i.e., between the ends of SIDs and the beginnings of STARs) appear to offer less clear-cut opportunities to save operating costs. Perhaps savings somewhat less than 1 percent of cruise fuel could be realized.

Potential savings from inefficient cruise altitudes are difficult to quantify because factors other than fuel economy often govern the choice of altitude. In the example considered, the most significant inefficiency that might be addressed by NASA DSTs was the fairly large interval between present Instrument Flight Rules altitudes above 29,000 feet.

We lack sufficient information to estimate soundly the costs of avoidable groundside delays. However, a crude estimate, which we believe to be conservative, suggests that these costs may exceed $100 million per year at busy at the 12 busiest CONUS terminals, for 1996 traffic levels.

Our sense of the relative impacts of the inefficiencies that our review identified may best be expressed by comparing costs of the potentially avoidable delays we identified for ground operations, climbs and descents, at busy terminals. These delay times are 1 minute for each taxi-out and taxi-in, 3 minutes for climb and 7 minutes for descent. Table 3-5 suggest pricing a 1-minute taxi delay at $33.00, a 1-minute delay in climb phase at $39.00 (the vector-out cost), and a 1-minute descent delay at $36.00.

With these prices, the per operation avoidable delays that we have identified cost $33.00 each for taxi-out and taxi-in, $117.00 for climb and $252.00 for descent.

Chapter 3
A Method for Cross-Comparable Assessment of Decision Support Technologies Benefits

OVERVIEW

The principal result of this project is the proffering of a method for making mutually consistent estimates of DST benefits, both singly and in various combinations. Figure 3-1 shows the method schematically.

Figure 3-1. Schematic Diagram of the Method

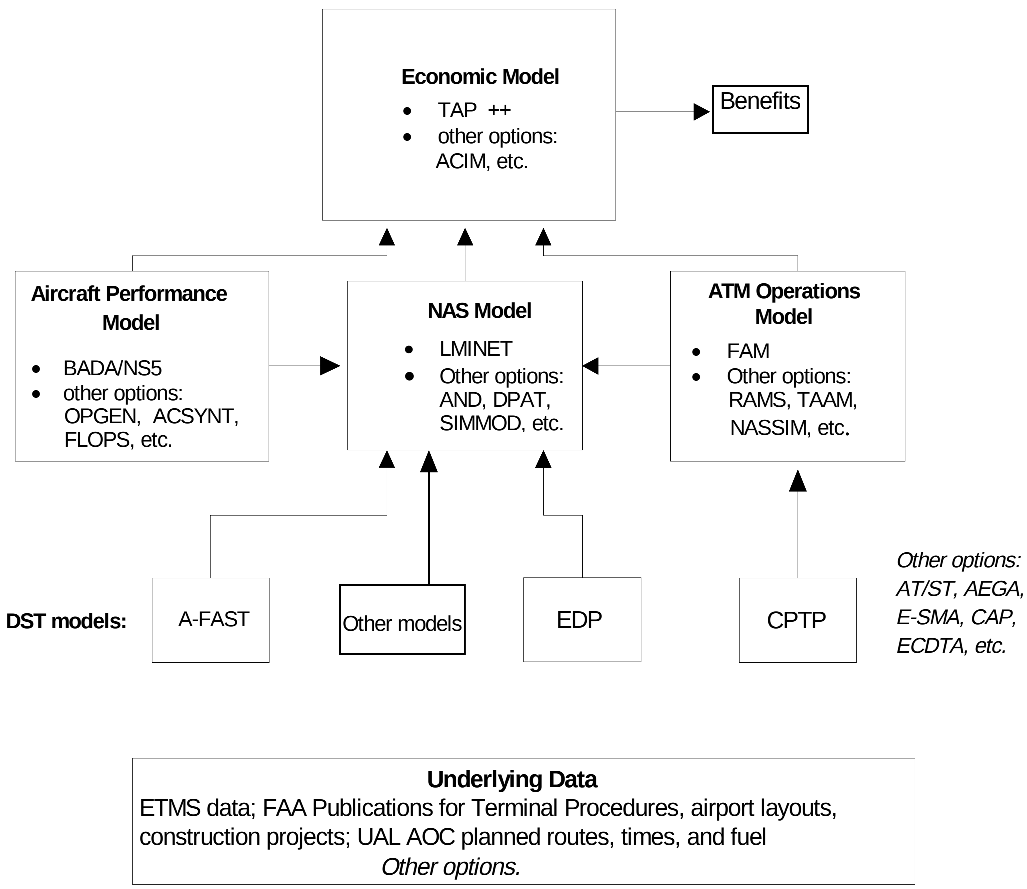

In Figure 3-1, the method's seven principal components (the capstone economic model, an aircraft performance model, National Airspace System model, an Air Traffic Management (ATM) model of operations, a set of DST models, and an underlying set of data are labeled in bold face type. Immediately following each of these labels, we show the particular instantiation used in the examples in this report. Following that, we list examples of other options that could be used to instantiate the component.

Applying the method is straightforward, though developing all the information required takes considerable effort. Applications go like this: for each DST to be assessed, the user determines the set of component models that will be used to reflect the tool's effects. For example, in our modeling of A-FAST, we noted that its effects could be reflected in changes to several parameters of LMINET's runway and airport capacity models. We did this fairly quickly, over a few days, because several of us had a good deal of current experience in modeling airport capacities.

In contrast to that rapid progress, when modeling CPTP, we found it necessary first to use FAM simulations to characterize the tool's effects and then to use the simulation results to determine changes in parameters of LMINET's en route sector models. Developing and operating the simulations took several weeks.

For modeling Expedite Departure Path (EDP), we considered FAM simulations, ETMS data, and the results of interviews with a controller who had experience in the NYC TRACON to arrive at an adjustment to certain of LMINET's departure TRACON parameters. Much of this work dovetailed with the CPTP work and our work to identify opportunities for NASA DSTs (see Chapter 2) Nevertheless, developing (with confidence!) the single parameter change for certain of LMINET's departure TRACON models took many hours' work by several people. EDP also furnishes an example of a DST whose benefits enter the economic models via two component models. EDPs expected reduction in times-to-cruise affects the mean time-in-sector of certain departure TRACON models in LMINET, and LMINET captures the resulting reductions in delays. The BADA aircraft performance models capture the airlines' savings in fuel, which turns out to be EDP's more significant effect.

Once models are determined for a set of DSTs, operating the NAS model gives before-and-after delay data to the economic model, and, as required, operating the aircraft performance model gives savings in block fuel and block time, while operating the ATM model may give information about savings resulting from improved controller productivity.

NATIONAL AIRSPACE MODELS

An essential feature of our assessment method is a model of the national airspace system. The model may be either a queuing network model, as implemented here, or a simulation model, as described briefly in Chapter 3. Whichever model is chosen, it must have sufficient granularity to capture the effects of the DSTs under review.

LMINET

LMINET is a queuing network model of the NAS developed by LMI for NASA [1]. In general terms, LMINET models flights among a set of airports by linking queuing network models of airports with sequences of queuing models of TRACON sectors and en route sectors. The user may specify the sequences of sectors to represent various operating modes for the NAS. The sequences may, for example, correspond to optimal routes for the winds aloft of a specific day, or they may correspond to trajectories of flights as flown on a specific day, as determined from ETMS data.

The network is driven by a schedule of departures from its airports and by a schedule of arrivals from outside the network. The Official Airline Guide is one source of such schedules of departures and arrivals. Both airport and sector capacities may be affected by weather. Weather data are provided to LMINET as epoch-by-epoch values of meteorological conditions at each of the airports and as epoch-by-epoch values of a single weather parameter for each TRACON and en route sector.

The following subsections describe more details on LMINET's components.

Airport Models

Each airport model is itself a queuing network, as shown in Figure 3-2.

Figure 3-2. Queues in the LMINET Airport Model

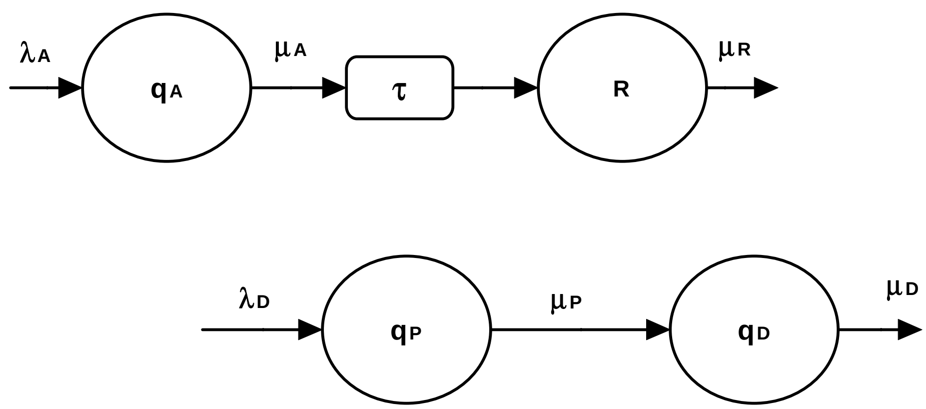

Traffic enters the arrival queue, q_A, according to a Poisson arrival process with parameter $\lambda_A(t)$. Upon service by the arrival server, which is Poisson service with parameter $\mu_A(t)$, and after the turnaround delay τ, arriving aircraft enter the ready-to-depart reservoir R. Each day's operations begin with a certain number of aircraft in this reservoir.

Departures require two services: an aircraft and a departure runway. Departures enter the queue for aircraft, q_P, according to a Poisson process with rate λ_D. Departure aircraft are assigned by a process with service rate $\mu_P(t)$. When a departure aircraft is assigned, R is reduced by one. Having secured a ready-to-depart aircraft, the departure leaves q_P and enters the queue for a departure runway, q_D, where it is served according to the departure process characterized by $\mu_D(t)$.

Service at the queue for departing aircraft depends on the state of the ready-to-depart reservoir R. If R is not empty, then the service rate $\mu_P(t)$ is very large compared with one (service time is very short). If R is empty, then departing aircraft are supplied by output of the arrival queue, delayed by the turnaround time τ. The service processes to R and q_P are given by

$$\text{Service to R:} \begin{cases} \lambda_D, R > 0 \\ o_A(t-\tau), \text{else} \end{cases}$$

and

$$\text{Service to } q_P: \begin{cases} M >> 1, R > 0 \\ o_A(t-\tau), \text{else} \end{cases}$$

where o_A is the output of the arrival queue.

Since aircraft are not interchangeable, this assumption on the supply of departing aircraft is tenable only when delays in the arrival process do not significantly alter the sequence of arrivals.

Service rates to the arrival and departure runways, λ_A and λ_D, respectively, are determined by individual airport capacity models that generate arrival and departure capacities as functions of meteorological conditions (e.g., ceiling, visibility, wind speeds and direction) and arrival and departure demand. Several parameters characterizing a specific airport affect the airport capacity models, as shown in Table 3-1.

Table 3-1. Runway Capacity Parameters

Symbol	Definition
c	Mean communication time delay
δc	Standard deviation of communication time delay
D	Length of common approach path
D_D	Distance-to-turn on departure
ρ_i	Fraction of operating aircraft that are type *i*
RA_i	Mean arrival runway occupancy time of *i*th aircraft type
δRA_i	Standard deviation of arrival runway occupancy time of *i*th aircraft type

Table 3-1. Runway Capacity Parameters (Continued)

Symbol	Definition
RD_i	Mean departure runway occupancy time of *i*th aircraft type
δRD_i	Standard deviation of departure runway occupancy time of *i*th aircraft type
Sij	Miles-in-trail separation minimum, aircraft of type *i* behind aircraft of type *j*
V_i	Approach speed of aircraft type *i*
δV_i	Standard deviation in approach speed of aircraft type *i* δ
δW_i	Wind variation experienced by aircraft of type *i*
δX_i	Standard deviation of controller's information on position of aircraft *i*

In addition to the runway capacity parameters, LMINET's airport capacity models respond to information on the configurations in which the airport is usually operated.

Presently, LMINET is implemented with 64 airports.[1] Figure 3-3 shows their locations. They account for over 80 percent of the air carrier operations for 1997, as reported in the current *FAA Terminal Area Forecast*. The LMINET airports are a superset of the FAA's 57 pacing airports.

Figure 3-3. LMINET Airports

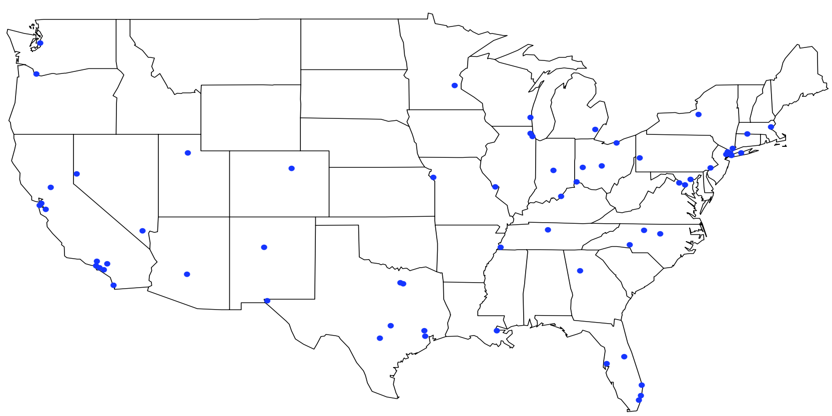

[1] The 64 airports are ABQ, ATL, AUS, BDL, BNA, BOS, BUR, BWI, CLE, CLT, CMH, CVG, DAL, DAY, DCA, DEN, DFW, DTW, ELP, EWR, FLL, GSO, HOU, HPN, IAD, IAH, IND, ISP, JFK, LAS, LAX, LGA, LGB, MCI, MCO, MDW, MEM, MIA, MKE, MSP, MSY, OAK, ONT, ORD, PBI, PDX, PHL, PHX, PIT, RDU, RNO, SAN, SAT, SDF, SEA, SFO, SJC, SLC, SMF, SNA, STL, SYR, TEB, and TPA.

Sector Models

Recent work at the Institute has produced new models of both ARTCC sectors and TRACON sectors as multiserver queues, specifically as $M/E_k/N/N+q$ queues. That is, as queues with Poisson arrivals, service times with the Erlang distribution with parameter *k*, and *N* servers. Not more than q clients will wait for service, so that the maximum number in the system is $N + q$.

The models were developed in light of several interactions with FAA people, including controllers at the Denver ARTCC and the Denver TRACON as well as experienced supervisory controllers working at the FAA's National Command Center in Herndon, Virginia. The development and calibration of the queuing models of sectors is described in Reference [1]. The following subsection gives some details of the model and our numerical treatment for operating LMINET.

THE $M/E_3/N/N+q$ SECTOR MODEL

In our queuing model for the ARTCC and TRACON sectors of the NAS, the times between aircraft arrivals to each sector are assumed to have the Poisson distribution. The time that an aircraft stays in a sector is assumed to be a random variable distributed according to Erlang-3 distribution. A sector can simultaneously handle no more than N aircraft at a time, when the capacity N is determined by the sector's characteristics and the weather. We also assume that, at most, q aircraft will "wait," (i.e., be delayed by speed changes or vectoring, to be served in a sector).

The arrival demand to a sector is determined by the network flight schedule. The choice of the Erlang-3 distribution for the times-in-sector was made in view of ETMS data and is explained in Reference [1]. We chose 18 as the maximum number of aircraft that a sector's controllers can handle at one time, to be consistent with Reference [11]. We base our choice of the maximum number of "waiting" aircraft on interviews with controllers at the Denver ARTCC.

Solving the model poses a great challenge to us, for there is no closed form solution, not even for the steady state for the $M/E_k/N/N+q$ queue. We have to resort to determining the probabilities of each state of the system numerically.

That is itself a respectable challenge, because the number of states is large. For a $M/E_3/N/N+3$ system, there are 1,950 states [2]. The number of states increases rapidly with N. For example, if q = 3 the number of states is 27,000 if N is 50; the number of states is 192,000 if N is 100; and the number of state is 620,000 if N is 150. Thus, determining the state probabilities directly from the evolution equations means solving a very large system of ordinary differential equations.

The systems' plant matrices are sparse, and the systems seem reasonably well-conditioned, so that brute-force numerical methods may succeed for some cases.

We have, in fact, generated numerical solutions of the full equations for N=18 and q = 3 in this way, to have means of checking the results of approximate solution methods. This approach takes too much time, however, to be at all appealing for routine use. Fast-executing approximate solutions are greatly desired. The trick lies in reducing the number of states.

Our key idea to improve the computer execution involves a new concept called mega state. The Erlang-3 distribution is equivalent mathematically to the distribution that results from service by three servers in tandem, each of which has the same Poisson distribution of service times. Thus, the state of a $M/E_3/N/N+q$ system is determined by four numbers i, j, k, and q, where i denotes the number of aircraft that have not completed one service of the three required, j denotes the number that have completed one but not two services, k is the number that have completed two but not three and q is the number of aircraft waiting.

The mega state m is defined as $m = i + j + k$. If the sector capacity is N, then $m \in [0,N]$. After checking the state transition matrix, we realized that a state interacts only with states of neighboring mega states. This further implies that for mega states m_1, m_2, $m_1 < m_2$, if $Pr(m_1)=0$ then $Pr(m_2)=0$, which can be proved by mathematical induction.

In practice, we can maintain a dynamic upper bound of mega state such that the probability of any mega state less than this upper bound is nonzero and the probability of any mega state equal or larger than this upper bound is negligibly small. Therefore, we do not need to solve all the state transition equations; we need to solve only the ones whose mega state is equal to or less than the upper bound. This technique alone reduces more than 90 percent of computer execution time. Since the upper bound is dynamic, there is virtually no loss of accuracy of solution, which we have verified by comparison with exact solutions.

For solving those state evolution equations that must be solved, we have tried forward Euler and second- and fourth-order Runge-Kutta integration schemes. Of the three, the second-order Runge-Kutta gives us the best speed, contrary to the conventional wisdom that the fourth-order Runge-Kutta would give the best speed. The higher the order in the Runge-Kutta integration scheme, the more accuracy we may get; hence, we may afford larger integration steps to speed up the process. However, due to the shear large number of differential equations that we have to deal with, some kind of stiffness must exist to prevent us from using large steps. We finally settled in the second-order Runge-Kutta scheme with the adaptive step. The adaptive step control works as follows. In moving the time by one step, we also move the time by two half steps. We then compare their results. If their difference is smaller than a specified number, we will enlarge the step in the next iteration; if their difference is larger than a specified number, we will reduce the step and go back to redo this integration step. Their difference is also used to get better precision. In working out several cases, we find that we gain a small fraction of the total time by using a second order Runge-Kutta scheme with adaptive step size.

Another important method for keeping the queuing calculations tractable is to introduce subsectors. This is particularly helpful for the rectangular-area sectors of LMINET, which can have large peak demands.

In operating the NAS, the FAA does subdivide busy sectors, geographically and/or by altitude. We model this in our subdividing busy sectors into sets of independent sectors, each of which has the N of a single sector. We have been careful not to carry this process beyond the point at which the subdivisions are at least arguably feasible for actual operations.

LMINET's rectangular en route sectors are roughly 120 miles on a side. They represent airspace above Flight Level 230. With present altitude-direction conventions, this affords about 14 levels at which modern turbojet transports may cruise: eastbound traffic at flight levels 230, 250, 270, 290, 330, 370, and 410; westbound traffic at flight levels 240, 260, 280, 310, 350, 390, and 430.

Thus, division into two subsectors can be accomplished feasibly either by altitude or geographic sectioning: two geographic subsectors would be 60 x 120 nautical miles, and two altitude subsectors would each have seven available flight levels.

Subsectoring with two geographic subsectors and two altitude subsectors is also feasible, so divisions with four subsectors is feasible.

Subsectoring into three geographic regions could certainly be accomplished feasibly, giving sectors 40 x 120 miles. Division of a rectangular sector into three subsectors by altitude division probably is feasible, as well: each subsector would have at least two altitudes. But the resulting combination, giving nine subsectors, may be about as far as one should go.

Internally, LMINET assumes that aircraft arriving at a subsectored sector are roughly evenly divided among the subsectors. Queue statistics are generated for just one of these, so, to get overall delay statistics, one scales up the single-sector result by the number of subsectors. The advantage for the queuing calculations is, that we never consider a sector capacity N larger than the value, typically 18, that is characteristic of a single controller team.

With mega states and subsectoring, and compiling the C code in which LMINET is written to optimize execution speed, we can generate statistics for one, 20-hour "day" of CONUS operations in roughly 15 minutes on LMI's HP D370 with RISC 2.0.

TRACON MODELS

Each airport's TRACON is modeled with two arrival sectors and one departure sector. The sectors are modeled as M/Ek/N/N+q queues.

LMINET allocates arrivals to an airport so that each arrival TRACON sector sees roughly half of the arrivals in each epoch of operation. For the work reported here, an epoch is 1 hour long.

EN ROUTE SECTOR MODELS

Like the TRACON sectors, en route sectors are modeled as M/Ek/N/N+q queues.

Automatic Traffic Flow Controller

This element of LMINET models the FAA's practice of delaying scheduled aircraft departures to congested airports. The function of this module is, essentially, to limit the arrivals to each airport by the airport's arrival capacity for each time epoch of the day, so that large arrival queues never form.

To perform this function, we construct a planning window, which is composed of the rest of day, to facilitate the planning of ground-hold decisions. At each epoch of the day, the module checks each airport's arrivals for the rest of the day. If the scheduled arrivals exceed the arrival capacity, the module will move some arrivals to the next epoch so that arrival demand meets capacity.

This process continues successively to the end of the day for each airport. Once this is done, the departure schedule is permanently changed based on the delays calculated during the process. The arrival queue and departure queue at the end of the last epoch are counted as additional demands to arrival and departure at the current epoch in the planning window, and the queue for planes from last epoch is counted as demand to both arrival and departure at the current epoch.

Even with the traffic flow controller, we cannot totally eliminate the arrival queues due to the facts that (1) we cannot delay an aircraft that is already in departure; (2) we will not delay the arrivals from the out-of-network airport; and (3) airport capacities are dynamic and dependent upon both arrivals and departures, which means that arrivals may exceed the arrival capacity even if arrivals equal capacity in the planning due to the large departure demand.

We implement the automatic flow controller by taking the following guidelines:

- Only departures *to* the congested airport will be delayed. The amount of delay is equally distributed among all the flights that are eligible to be delayed. We will not delay the departures *from* the congested airport to reduce congestion.

- Only the flights in the network airports may be delayed. The departures from airports out of the 64 airport network will not be delayed.

- We assume each airport is independent in its traffic flow control planning, and the decision to delay flights to the congested airport is solely based on the current schedule, current delays and queues and forecasted airport capacities. Since the air traffic flow control planning is done at each epoch for the rest of the day for each airport, the network effect of the traffic flow control is done through the modified schedule for the rest of the day. TRACON congestion is not a decision criterion.

- Local weather information, for the rest of the day, is assumed to be known to the air traffic controller at any time of the day.

- A flight can be delayed repeatedly as long as it has not yet departed.

The typical cause of airport and TRACON congestion is inclement weather, which will reduce both capacities. However, as we found out, we do not need to specifically count TRACON congestion as decision criterion, since once the arrival and departures are curtailed, the demand to the associated TRACONs will also be reduced.

Adjusting LMINET to Model the NAS in 1996 and 2005

Users may adjust several LMINET inputs: demand profiles, airport capacity models, sector capacity models, surface weather and weather aloft, routes between airports, and so on. This subsection explains our choices for the present instantiation of our DST assessment method.

DEMAND MODELS

We used the operations forecasts given in the electronic reference, "1996 Aviation Capacity Enhancement Plan and Airport Database," distributed as digital data on a compact disc by the FAA's Office of System Capacity [3] to model future demand for the airports of the 64-node LMINET model. The FAA's percent age growth in operations value is for the 15-year period of 1995 to 2010. Those forecasts, and the annualized growth rates, are shown in Table 3-2.

We developed demand inputs to LMINET for 2005, by scaling up departures at each network airport by the annual rate shown in Table 3-2, compounded for 9 years. We also increased the out-of-network arrivals at each airport by the same factor as the one used to scale up departures. We scaled up out-of-network arrivals to the sectors by 24 percent, as representative of the overall traffic growth expected from 1996 to 2005.

Table 3-2. Demand Growth Rates

Airport	FAA forecast 15-year growth in operations (percent)	Annualized growth rate (percent)	Airport	FAA forecast 15-year growth in operations (percent)	Annualized growth rate (percent)
ABQ	24	1.4	LGB	N/A	0
ATL	41.3	2.3	MCI	42.6	2.4
AUS	29.6	1.7	MCO	54.8	3
BDL	15.7	1	MDW	22.5	1.4
BNA	27.3	1.6	MEM	49	2.7
BOS	12.5	0.8	MIA	61.3	3.2
BUR	49.2	2.7	MKE	31.5	1.8
BWI	27.3	1.6	MSP	33.2	1.9
CLE	47	2.6	MSY	8.2	0.5
CLT	29.9	1.8	OAK	13.9	0.9
CMH	12.2	0.8	ONT	26.3	1.6
CVG	47.7	2.6	ORD	30.9	1.8
DAL	5.9	0.4	PBI	10.7	0.7
DAY	19.7	1.2	PDX	34.2	2
DCA	1.5	0.1	PHL	13.4	0.8
DEN	22.7	1.4	PHX	40.8	2.3
DFW	39.8	2.3	PIT	18.8	1.2
DTW	35.3	2	RDU	26.2	1.6
ELP	18.5	1.1	RNO	61.6	3.3
EWR	22.5	1.4	SAN	34.2	2
FLL	20.1	1.2	SAT	26.7	1.6
GSO	28.7	1.7	SDF	25.4	1.5
HOU	20.1	1.2	SEA	38.2	2.2
HPN	N/A	0	SFO	31.8	1.9
IAD	34.3	2	SJC	13.5	0.8
IAH	53.8	2.9	SLC	42.7	2.4
IND	35.6	2.1	SMF	59.3	3.2
ISP	-4.9	-0.3	SNA	23.8	1.4
JFK	17	1.1	STL	25	1.5
LAS	34.2	2	SYR	8.4	0.5
LAX	37.8	2.2	TEB	N/A	0
LGA	7.8	0.5	TPA	46.8	2.6

Some Important Caveats

To develop 2005 demands in this way is to assume that departures from a given airport will increase in the same ratio for all destinations. This is not likely to happen, but we have no satisfactory way to predict how the distributions of departures will change. To develop such predictions is beyond the resources of this task.

Also—and, we believe, quite significantly—air carriers are not likely to retain present schedules if doing so would result in serious delays. Faced with substantial delays, carriers probably would exercise such options as opening new hubs, operating more city-to-city services and fewer hub-and-spoke routes, and/or changing schedules to smooth out peaks in scheduled departure rates. Here again, adequately modeling carriers' responses to significant changes in the NAS seems to us to require a substantial effort that could not be accommodated in the present task.

Our method for DST assessment can readily account for actual changes in departure distributions and carriers' policies, when they are available. For now, the reader should bear in mind the limitations of our present demand model.

Capacity Models

We developed capacity models for the 64 LMINET airports for two periods: 1996 and 2005. The 1996 model serves as a basic reference and for comparisons with data on present-day NAS operations. The 2005 model is the reference for assessing DST effects. It includes planned FAA upgrades. We reviewed the FAA's *1996 Aviation Capacity Enhancement Plan and Airport Database* and *National Airspace System Architecture Version 2.0*, (FAA Office of System Architecture and Program Evaluation [ASD], October 1996) to determine these.

The following subsections describe our capacity models for 1996 and 2005.

Airport Models

We made a set of 32 airport capacity models to characterize the capacities of the 64 LMINET airports. There are specific models for several airports, including CLT, MDW, IND, JFK, EWR, DTW, SFO, ORD, DFW, and LAX. Other airports' capacities could be described by one of a group of standard capacity models. For example, many smaller airports can be characterized as providing two independent runways in VMC and one runway in IMC, whatever the wind.

Development of all the models was guided by the same considerations that govern LMI's capacity modeling (described in *Estimating the Effects of the Terminal Area Productivity Program,* NASA Contractor Report 201682, April 1997) necessarily simplified by the time and manpower limitations of the present study.

For the 2005 baseline, we included only those few airport construction projects described in the ACE database that would (1) be finished after 1996 but before 2005 (2) clearly increase capacity and (3) had *approved* environmental impact statements. These projects are as follows:

- DEN: Runway 16R/34L
- DTW: Fourth north-south parallel runway, Runway 4/22

- LAS: Upgrade of Runway 1L/19R to accommodate air carrier traffic
- MEM: New north-south paralleled Runway 18L/36R
- PHL: Commuter runway, Runway 8/26
- SDF: Replace Runway 1/19 with two new parallel runways separated by 4,950 feet, Runways 17R/35L and 17L/35R
- LAX: Remove 84/hour arrival-rate maximum imposed by groundside capacity limits.

Runway Capacities

Our airport capacity models require as inputs four parameters characterizing runway capacities as Pareto frontiers, for each of five meteorological conditions (VMC1, in which IFR flights may be concluded by VFR approaches; minimal VMC, the standard 1,000-foot ceiling and 3-mile visibility; ILS Category I; ILS Category II; and ILS Category III). We built these using the LMI Runway Capacity Model (described in *Estimating the Effects of the Terminal Area Productivity Program*, NASA Contractor Report 201682, April 1997).

We developed Pareto frontiers for two mixes of aircraft types. We refer to the two mixes as “domestic” and “international.” They characterize airports with mostly domestic traffic and airports with significant international traffic, respectively. The domestic mix is 10 percent small, 80 percent large, and 5 percent each for B757 and large; the international mix is 10 percent small, 60 percent large, 10 percent B757, and 20 percent heavy.

Our review of National Airspace System Architecture Version 2.0 led us to conclude that the chief capacity-enhancing improvements planned by the FAA that will be in effect by 2005 are CTAS build 1 and build 2. These two CTAS builds, which will include P-FAST, are to equip only eight airports, which are not specifically identified.

Lacking information about the specific airports that the FAA will equip, we decided to include CTAS at 10 key airports: BOS, EWR, JFK, LGA, ATL, DFW, LAX, SFO, ORD, and DTW, for our 2005 baseline. Because the specific effects of P-FAST are so closely related to those of A-FAST, they are described below in the section, “DST Models,” page 3-22.

We also adjusted the runway capacity models in LMINET to account for current departure visibility rules negotiated between the FAA and most major airlines, for many of the LMINET airports. As given in the several volumes of *U. S. Terminal Procedures*, (U. S. Department of Commerce, National Oceanic and Atmospheric Administration, National Ocean Service, various dates), the standard IFR takeoff minimums are one statute mile for aircraft with two or fewer engines, and one-half statute mile for aircraft with more than two engines. However, according to information that the FAA makes available at http://www.faa.gov/avr/afs/afs410/Cat3_ac.txt, most major U. S. airlines have negotiated operating rules for many airports, that enable them to take off with runway visible ranges of 700 feet or even less. In view of this information, we adjusted the LMINET capacity models to permit departures when visibility was 700 feet or more, at ATL, BDL, BNA, BOS, BWI, CLE, CLT, CMH, CVG, DAY, DCA, DEN, DFW, DTW, EWR, IAD, IAH, IND, JFK, LAX, LGA, MCI, MCO, MEM, MKE, MSP, MSY, OAK, ONT, ORD, PDX, PHL, PIT, SEA, SFO, SLC, SMF, and STL.

WEATHER MODELS

This subsection describes the way we modeled weather effects on operations of the NAS.

Data Sources

We obtained weather data from two principal sources. These are described in the following subsections.

Surface Weather

We obtained hourly reports of surface weather for 538 stations in the contiguous United States, from the National Climatic Data Center's On-Line Access and Service Information System, OASIS, for April 8, 1996, and June 12, 1996. We purchased data for November 29, 1996 from the National Climatic Data Center. With few exceptions, this provided hourly reports at each LMINET airport for all the days we considered.

The exceptions were of two kinds: a few missing reports in otherwise complete records, and the complete absence of records from some network airports. When only a few reports were missing in a record, we filled the gaps by interpolation.

When our OASIS searches gave no records for an LMINET airport, or when in our judgment gaps were too great to fill by interpolation, we used reports from the closest station with similar climate for which data were available. For example, for April 8, 1996, we used reports from EWR at TEB, and reports from APA (Denver Centennial Airport) at DEN.

We used data from LAX for ONT, even though other airports were slightly closer, because the FAA's reports of VFR days-per-year at those two airports were more alike than were those reports for the geographically closer stations.

The Days Considered

We used 3 days in 1996 as representative of typical weather patterns over the CONUS: April 8, June 12, and November 29. June 12 is our model of a “generally good” weather day. Except for brief periods of ILS Category I conditions, all 64 LMINET airports were VMC all day, with the exception of JFK, ISP (Islip, New York), and HPN (White Plains, New York.) Those three airports experienced mostly IMC. The busiest one, JFK, was in ILS Category I during its busy period, early afternoon through early evening.

April 8, 1996 is our model of a “somewhat degraded” weather day. It was characterized by generally good weather over most of the CONUS, with periods of IMC at certain terminals. A cold front running from Seattle to just north of San Francisco brought showers and reduced visibility to the northwestern corner of the CONUS. Another cold front pushing southward across the Appalachians produced showers and snow flurries in the northeast, notably at Boston, and into the mid-Atlantic states. A stationary front caused showers and some reduced visibility over the Florida peninsula. While there were some locally significant delays as at Boston, air traffic over the bulk of the CONUS seems not to have been significantly disrupted.

We used November 29, 1996 as an example of a bad weather day. There was rain in the far northwest and far northeast corners of the CONUS, and a storm spread snow over Colorado, western Iowa, New Mexico, northern Texas, and the Oklahoma panhandle causing periods of IMC at terminals in those areas. Chicago had IMC from the late afternoon through the end of the day. Dallas’ weather was IMC throughout the evening. Minneapolis-Saint Paul and Milwaukee were in IMC for much of the day. Kansas City experienced IMC all day.

Winds Aloft

We obtained values for eastward and northward components of winds aloft from NOAA’s National Center for Environmental Prediction (NCEP.) The data were created in the Atmospheric Chemistry and Dynamics Branch of NASA/GSFC. These were given on the same grid as the one described in Reference [1], and we used the same interpolating scheme as the one described there to generate smoothed wind fields.

Modeling Annual Variations with Representative Days

To make an estimate of yearly delay costs, we made weighted averages of delay costs for the three representative days: April 8, June 12, and November 29. A study of 30 years’ weather data at 10 busy airports that the Institute made for NASA in 1995 showed that the airports experienced VMC more than 80 percent of the time. They had ILS Category I conditions roughly 10 percent of the time.

Guided by this and by estimates of the costs of arrival delays at BOS made in Reference [14], we chose weights of 0.8 for the June day, 0.13 for the April day, and 0.07 for the November day.

Other Options for the Weather Model

The present instantiation generates statistics for DSTs' benefits by introducing actual CONUS-wide weather data into models that respond to them. An alternative would be to develop models of key weather parameters that appropriately reflected correlations among weather parameters at various sites. Such models could be developed from available weather data and then used with network or simulation models of the national airspace system.

Consistency of LMINET Results with OAG and ETMS Data

LMINET generates a good deal of information.

- statistics on traffic, on queues for arrival and departure service, and for departing aircraft, hour-by-hour at 64 airports; hourly statistics on traffic and delays at 64 departure TRACONS and 128 arrival TRACONs; and,
- with the sectorization that we used, hourly statistics on traffic and delays at 126 geographic-area en route sectors.

We found it helpful and encouraging to compare arrival and departure demands from LMINET, with demands from the OAG, and from ETMS data. Figure 3-4 shows results for DFW as an example.

Figure 3-4. Arrival and departure demands at DFW, from OAG, ETMS, and LMINET, for April 8, 1996

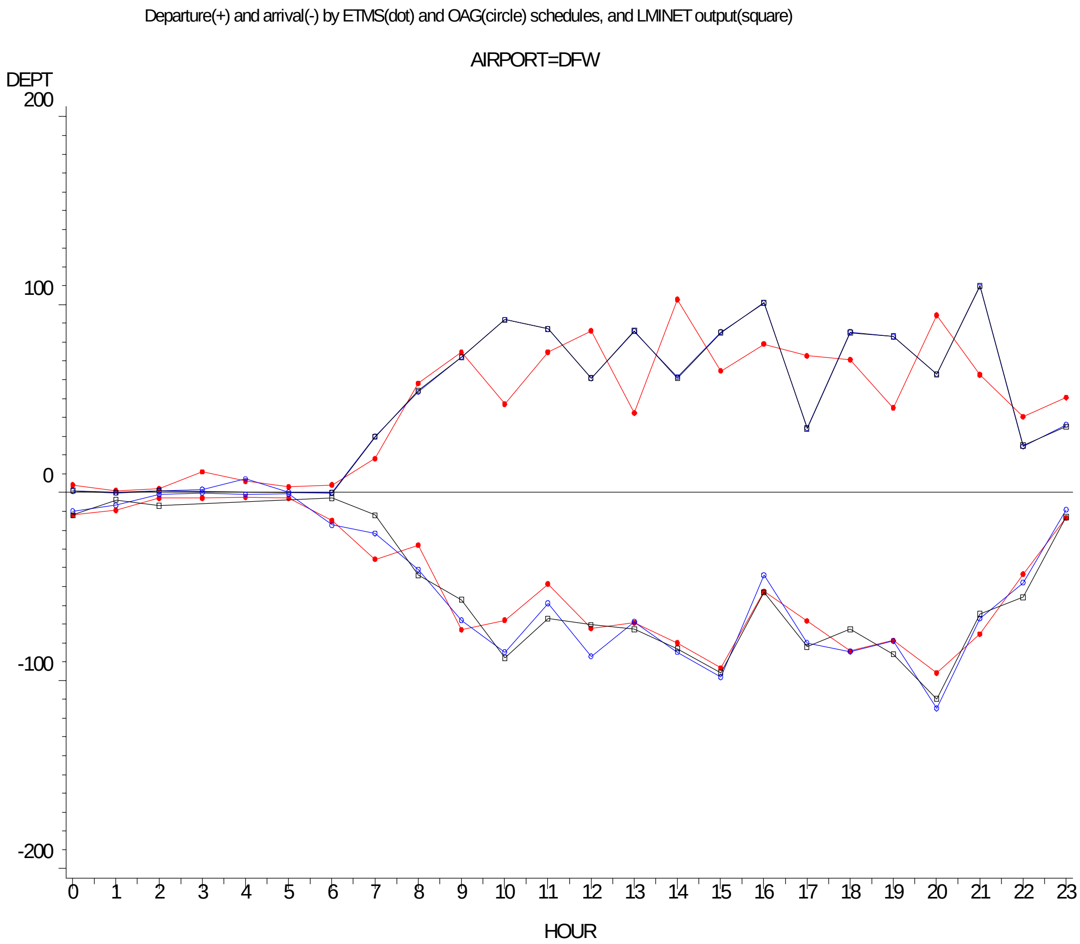

Since LMINET departures are derived from OAG data, those two traces agree quite closely. LMINET and OAG arrival demand data also agree reasonably well, and this fact provides a consistency check on LMINET's results. ETMS data do not agree particularly well with the OAG departure data, except for the overall scale. Interestingly, LMINET arrival data agree rather better with OAG data than do arrival data.

As shown in Figure 3-5, a similar situation obtains at ORD.

Figure 3-5. Arrival and departure demands at ORD, from OAG, ETMS, and LMINET, for April 8, 1996.

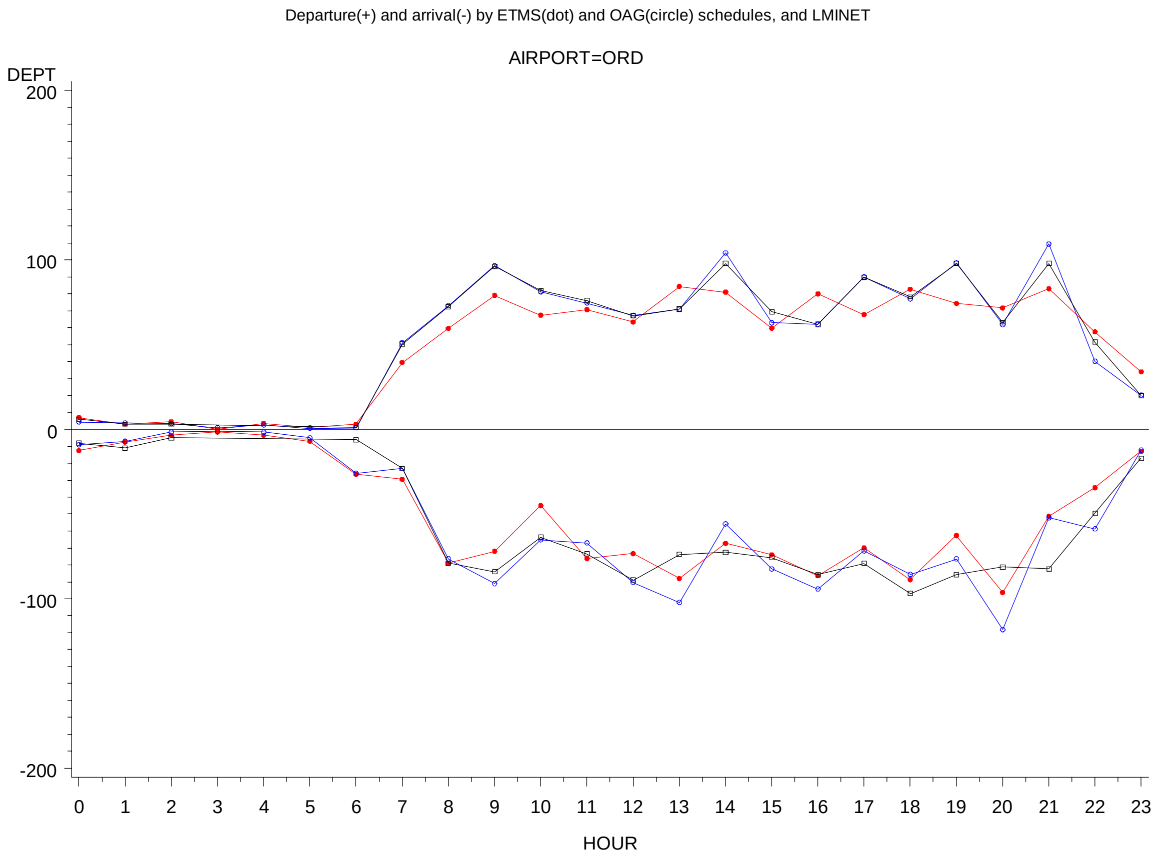

Other Options for the NAS Model

Several NAS models would be viable options for the NAS model, each bringing its own set of strong points and not-so-strong points. The MIT/MITRE Approximate Network Delays (AND) model is another queuing network model of the NAS. AND offers an itinerary builder which constructs feasible aircraft itineraries from OAG data, so that cumulative delay effects may be studied. Presently AND is being extended to include the same queuing model of en route and TRACON sectors that is in LMINET.

Detailed Policy Analysis Tool (DPAT) is a simulation model of the NAS, derived from the FAA's NASPAC simulation model. Until recently, conventional wisdom held that simulation models, while offering excellent opportunities for detailed models, took too long to set up and operate to be useful in situations where quick results were desirable. With Detailed Policy Analysis Tool (DPAT), however, MITRE staff have developed means of executing a simulation model sufficiently rapidly that it may well be useful for evaluating DSTs.

A simulation model like SIMMOD, which may not execute quickly enough when modeling substantial parts of the NAS to be useful in our method for assessing DSTs, may nevertheless be quite helpful if it is used to model only a small part of the NAS, such as a single important airport or small set of airports. In that case, it may be possible to exchange data between SIMMOD and a queuing network model like LMINET or AND to get whole-NAS results corresponding to quite detailed DST effects at specific airports.

ATM Operations Models

A model capable of representing effects of DSTs on operations of ARTCC sectors and TRACONs is a key part of our evaluation method. This section describes the model that we used, the Aircraft/Air Traffic Management Functional Analysis Model (FAM), in some detail and mentions some other options.

Functional Analysis Model

The Logistics Management Institute developed the prototype Aircraft/Air Traffic Management Functional Analysis Model (FAM) under contract with the NASA Ames Research Center [5]. FAM is a discrete event simulation model that runs on either the IBM/Windows'95 or Macintosh personal computers. This new model addresses the shortcomings of earlier simulation models by discretely modeling all of the active participants in the airspace system, airline air crews, as well as airline operations centers (AOCs) and FAA controllers.

To measure impacts of technologies on the ATM system, FAM accumulates statistics of interest on task loading and usage time for humans and equipment in aircraft, FAA ATC facilities, and AOCs. As a discrete event simulation, FAM carries out each single event in a series that occurs over time.

Figure 3-6 shows FAM's basic architecture. FAM generates events representing an interaction among components of the model. The events are passed between the affected objects in the model using the appropriate communications channels. For example, FAM currently models both sides of a conversation with a sequence of individual events. Each event represents one party speaking once. The initial, or "trigger" event in each sequence is released by the External Event Generator and sent to the model object that originates that event. For example, the opening transmission in a radio conversation would be sent to the model object that makes the transmission. From there, FAM routes the event through the appropriate channel in the Communications module to the destination object, the recipient of the transmission.

Figure 3-6. FAM Architecture

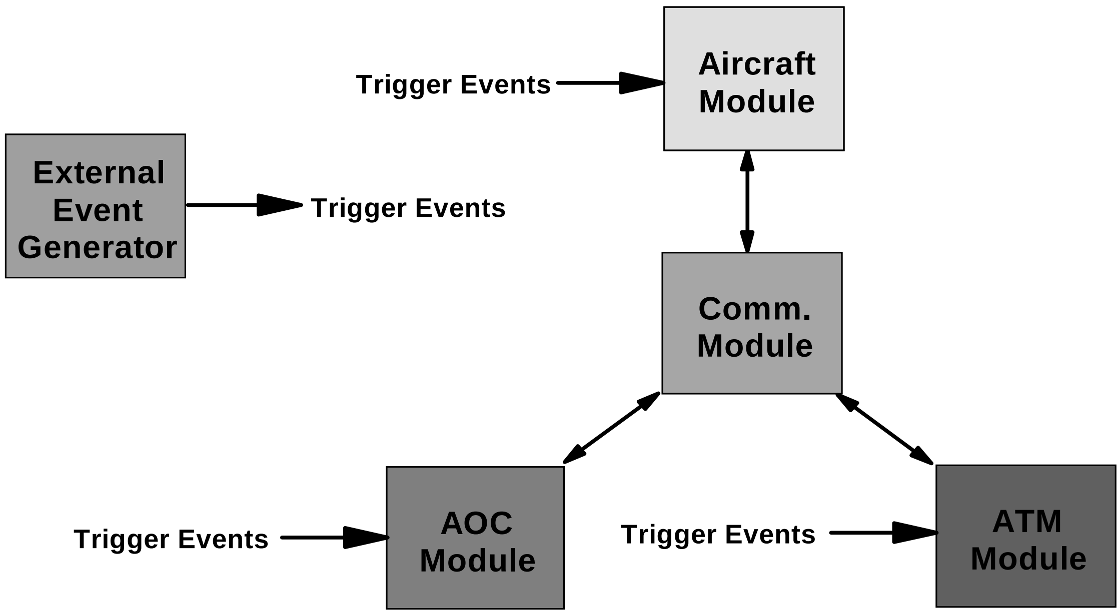

Each event carries with it a set of attributes, such as the event's origin and destination objects or the type of event. The model objects "take" appropriate action on the events, in many cases generating a reply event, on the basis of these attributes. Each model object accumulates workloads and equipment usage times associated with each event. These, together with time series of the simulation-sample means of such features as the number of aircraft in the models, are available to users.

Other Options for ATM Operations Models

The Total Airspace and Airport Modeller (TAAM) provides detailed simulations of ground and terminal area operations, and is capable of modeling en-route traffic as well. TAAM is a commercial product. The Reorganized ATC Mathematical Simulator (RAMS) is a simulation tool developed by Eurocontrol. Information about RAMS is available on the Internet, at www.eurocontrol.fr/projects/rams.

The National Airspace System Simulation (NASSIM) model is an FAA model, developed to simulate effects of new procedures and equipment on the NAS, for current and forecast air traffic levels[2].

AIRCRAFT PERFORMANCE MODELS

Aircraft performance models give fuel burns and flight times, and, with information on winds aloft, determine optimal trajectories. This section briefly describes the model used in our instantiation, and it mentions some other options.

The Base of Aircraft Data/Flight Segment Cost Model

The performance model that we used is the same one used in Reference [1]. The model has been developed into a part of a tool for predicting costs of operating a specified set of flights, suitable for general use and available on the World-Wide Web [6]. This latter use, in the Flight Segment Cost Model, generated the acronym FSCM.

As used here, the model generates fuel burns and flight times for aircraft to fly through a specified set of three-dimensional position points, at a prescribed set of Mach numbers, for given winds aloft. That is, given $\{latitude_i, longitude_i, altitude_i\}_0^{N+1}$ and $\{\mathrm{M}_\mathrm{i}\}_0^{\mathrm{N}}$, and also a function returning the vector-valued wind at a given latitude, altitude, and longitude, the model computes the fuel burn and flight time for a specified aircraft to fly the prescribed trajectory.

The model uses the flight mechanics equations and parameters of the Base of Aircraft Data models developed by the European Organization for the Safety of Air Navigation [7]. Available data cover more than 125 aircraft types. As detailed in Reference [1], fuel burns are computed from an exact solution of the BADA equations for flight at constant altitude and Mach number, corrected for climbs, descents, and speed changes by the total energy method.

Our input winds aloft are gridded data for eastward and northward winds aloft, provided generously and promptly by NOAA’s National Center for Environmental Prediction (NCEP.) The data were created in the Atmospheric Chemistry and Dynamics Branch of NASA/GSFC. We produce smooth values of winds aloft by the interpolation scheme described in Reference [1].

The model’s outputs agree well with others. Figure 3-7 shows the agreement between its predictions of fuel burns, and those of United Airlines’ air operations center, for several Boeing 727 flights in early August 1997. compares predicted

[2] Salanski, M., “NASSIM Simulation Development Overview,” Federal Aviation Administration, AOR-200, Operations Research Service, June, 1994.

block times. The smaller values of time and fuel are for trips from DEN to ORD, the larger for flights from ORD to PHX.

Figure 3-7. Comparison of BADA/FSCM and UAL AOC Predictions for Block Fuel

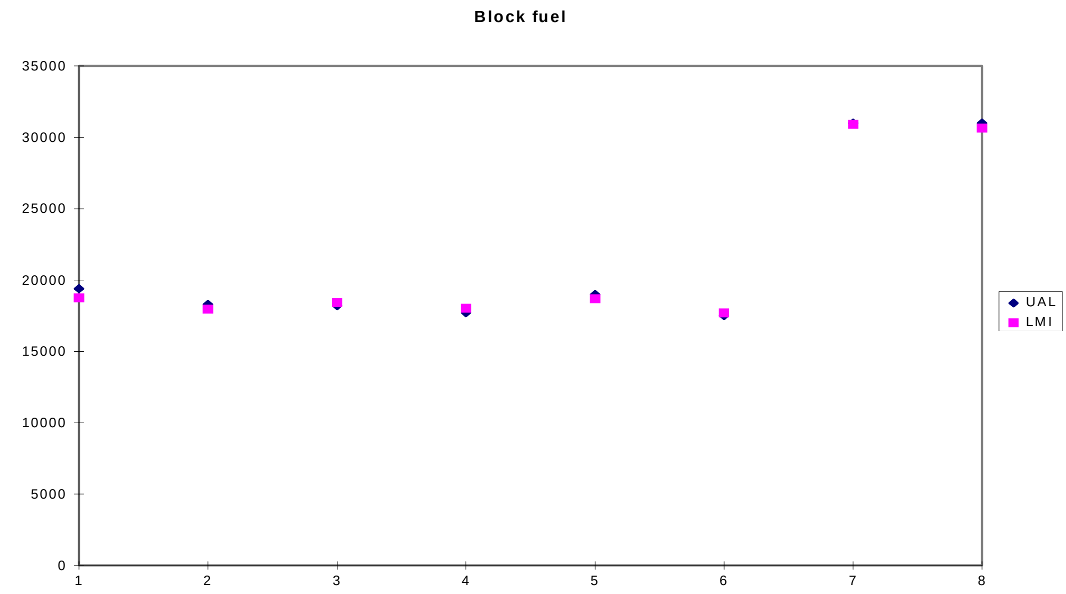

Figure 3-8. Comparison of BADA/FSCM and UAL AOC Predictions for Block Time

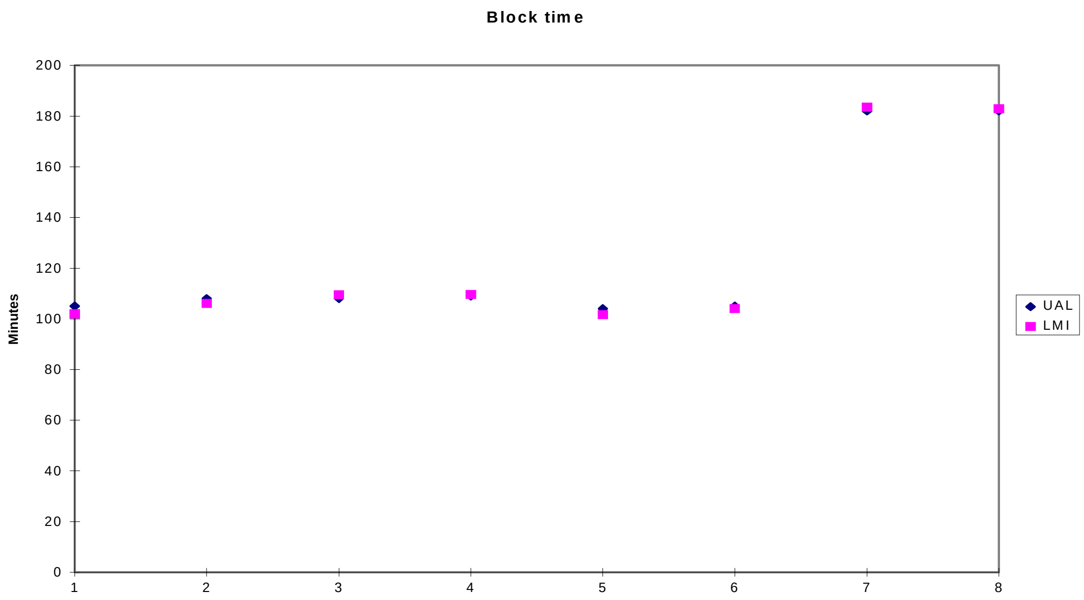

Other Aircraft Performance Models

Several other aircraft performance models could be used. If characteristics of new aircraft, or aircraft not in the BADA database, were important, then such NASA preliminary design models as ACSYNT [8] or FLOPS [9] could be used. The preliminary design models can be used directly to produce fuel burns and flight times, and they may be used with optimization routines to generate optimal routes. Alternatively, outputs from the preliminary design models may be used to infer appropriate BADA parameters for new aircraft; with these, the BADA/FSCM models can be applied.

DST MODELS

This section describes ways to model NASA DSTs with the instantiations of the component models used in the present study. The component models required are a NAS model, an aircraft performance model, an ATM operations model, and appropriate data. For this study, the instantiation of the NAS model is LMINET, the aircraft performance model is BADA/FSCM, and the ATM operations model is FAM.

The discussions in this section are inclusive, describing how specific DSTs can be brought into the present evaluation framework even when to do so requires more resources than the present study affords. In Chapter 4, we use available resources to model three DSTs in sufficient detail to generate preliminary benefits estimates.

The following subsections discuss general considerations for modeling DSTs and give some specifics for modeling a set of DSTs.

General Considerations for Modeling DSTs

LMINET may be adjusted at several levels, using any of the parameters of its constituent models, to reflect DST performance. At the highest level, airport capacities may be adjusted simply by multiplicative factors applied to arrival and/or departure capacities. At the most detailed level, DST effects may be reflected in changes to the runway capacity model parameters shown in Table 3-1.

The effects of DSTs on airspace outside airports may enter LMINET by adjustments to the parameters of the queues that model TRACON and en route sectors. A DST that reduced a controller’s workload might, for example, be reflected in an increase to the maximum number N of aircraft that could be accommodated at one time. A DST that, like EDP, reduced the amount of time aircraft spend in a sector as well as the controller’s workload, could be modeled by an increase in N and a decrease in the mean of the Erlang distribution of service times.

To develop values for the variations in LMINET parameters that model DSTs’ effects on sectors, we used the FAM simulation model.

Conflict Prediction and Trial Planning Tool

The Conflict Prediction and Trial Planning Tool (CPTP) will help en route sector controllers identify and resolve potential conflicts. Intended as a precursor of the Airspace Tool/Sector Tool DSTs described in Chapter 3, CPTP will serve as a research tool for developing those DSTs while assisting controllers.

CPTP will receive radar track and flight plan information from the host system and winds aloft data from the National Weather Service's Rapid Update Cycle predictions. These data, with extensions of CTAS' trajectory synthesis algorithms, will provide predictions of potential conflicts considerably in advance of those developed now by individual controllers.

CPTP will send warnings of identified potential conflicts to the display(s) of the controller(s) whose sectors are affected. Controllers may then use the "trial planning" feature of CPTP to test resolution strategies, before issuing clearances to the aircraft involved. For controllers directing aircraft in transitions between en route and terminal airspace, CPTP's trial planning functions have the ability to respect any imposed miles-in-trail restrictions.

Models of CPTP must capture the tool's effects on individual sector operations and on the NAS as a whole. The latter task can be done by a queuing network model such as LMINET. Such models characterize sector performance by only a few parameters: LMINET uses just three, namely, the maximum number of aircraft that a controller team can handle at one time in a given sector, the index k of the E_k distribution of times-in-sector that characterizes the degree to which times-in-sector are concentrated about their mean, and the mean time-in-sector.

Detailed analyses of sector operations are required to generate numerical values that characterize the changes CPTP may be expected to make in the sector model's parameters. In the present work, we used the NASA Functional Analysis Model. Developed for NASA by LMI, FAM is a discrete-event simulation model capable of modeling sector operations in considerable detail.

We set up FAM to model one sector in the Denver Air Route Traffic Control Center (ZDV), together with the Denver TRACON and the Denver and Colorado Springs airports. For each of the parts modeled, FAM monitors the utilization of the controllers and operators. We took the basic demand event file that FAM uses for this task directly from actual ETMS data, for flights that flew through the sector that we considered. To increase demand in the sector, we modified the original event file. The model simulates a 4-hour period of operations.

We derived the initial conflict resolution time of 50 seconds used in the model from an average of the 40 and 60 seconds that Grossberg, Richards, and Robertson report it takes to resolve "crossing conflicts and overtaking conflicts, respectively." [10] This is a fixed conflict resolution time, for the purposes of this model. In our simulation, conflicts are generated by a random event generator that produces events based on the number of aircraft in the sector.

When we modified the original event file to create a change in the maximum and total number of aircraft in the sector, we maintained the same distribution of demand. The following figures (3-9 to 3-16) show the demand distributions for the event files used in the analysis. Notice how the distributions for the demand stay fairly consistent.

Figure 3-9. Demand Distribution for Utilization 65.3 percent

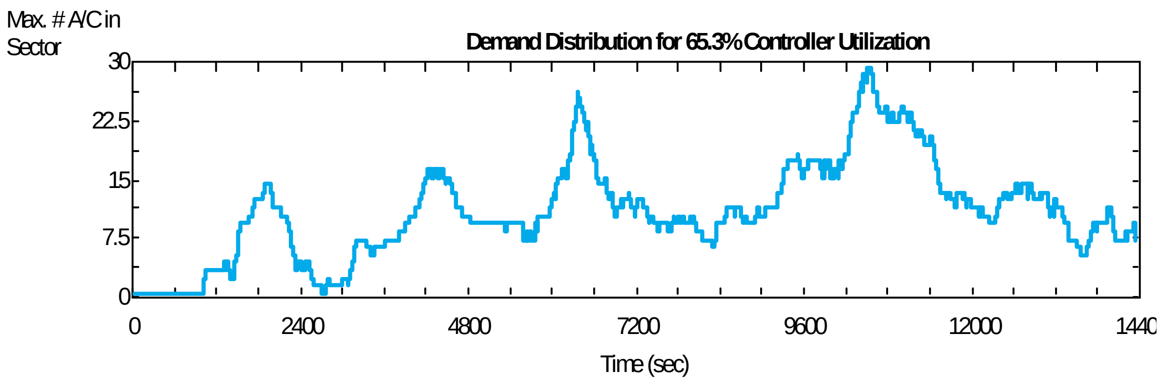

Figure 3-10. Demand Distribution for Utilization 63.74 percent

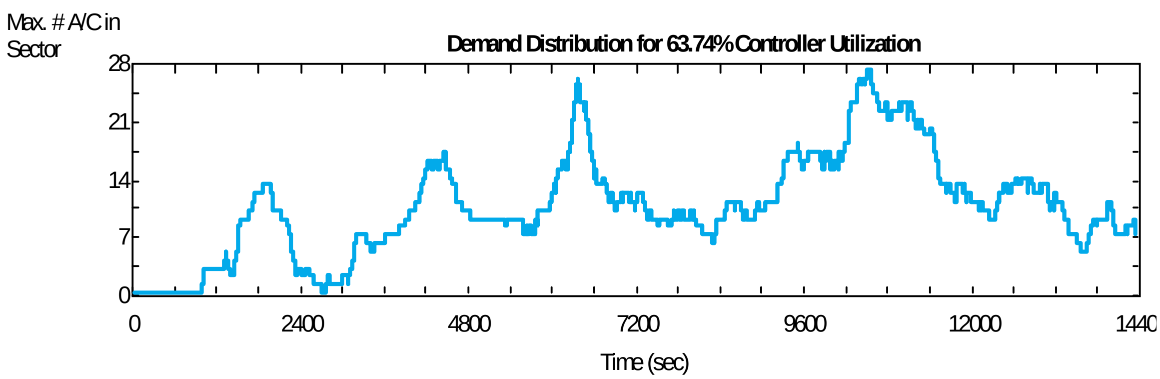

Figure 3-11. Demand Distribution for Utilization 61.79 percent

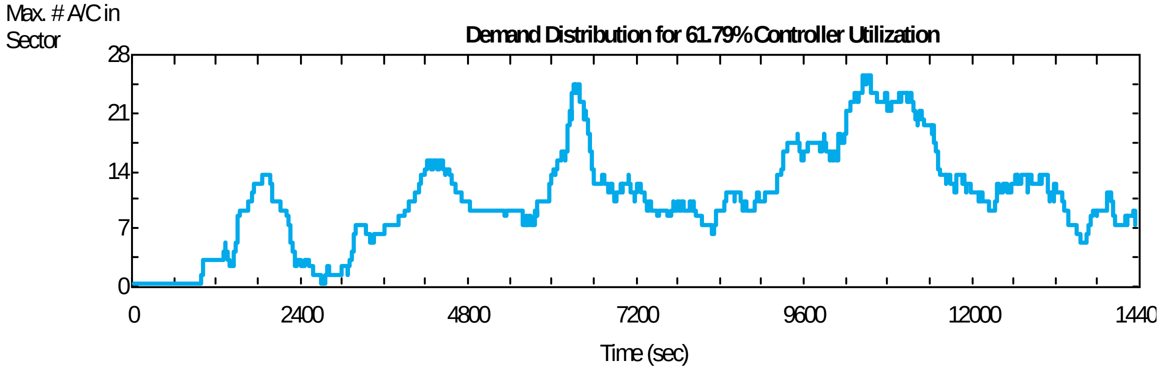

Figure 3-12. Demand Distribution for Utilization 59.46 percent

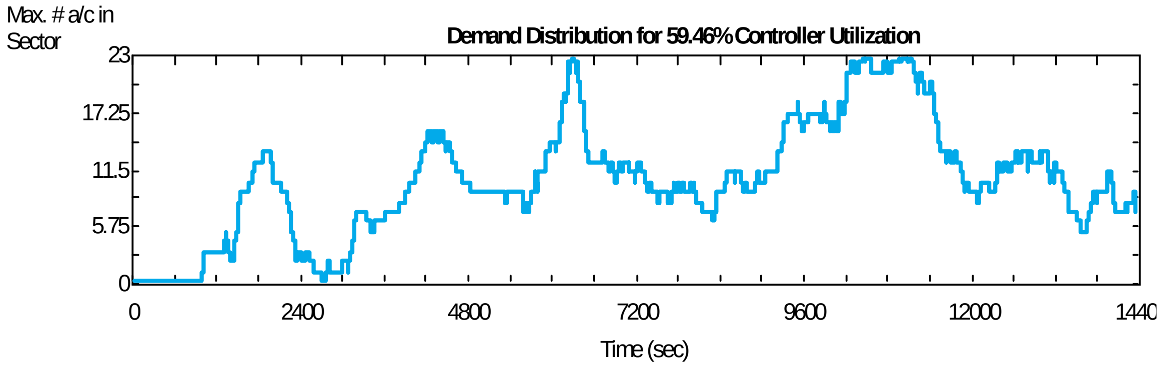

Figure 3-13. Demand Distribution for Utilization 57.90 percent

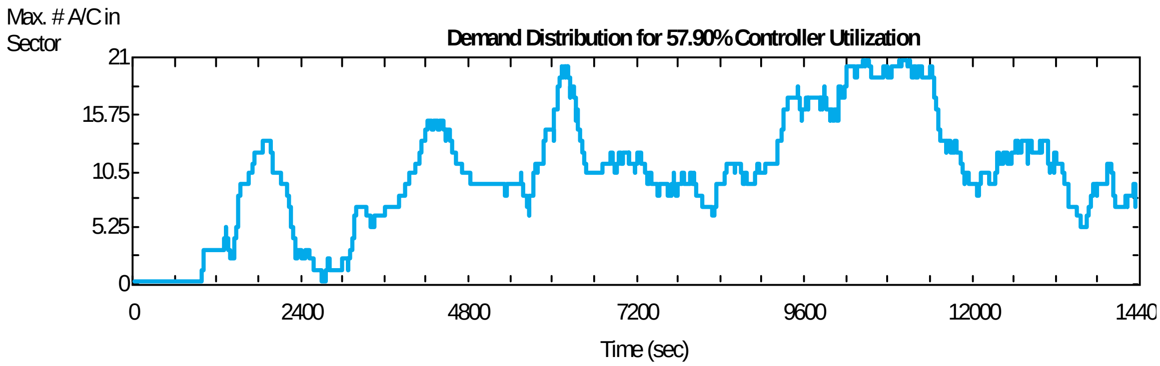

Figure 3-14. Demand Distribution for Utilization 55.94 percent

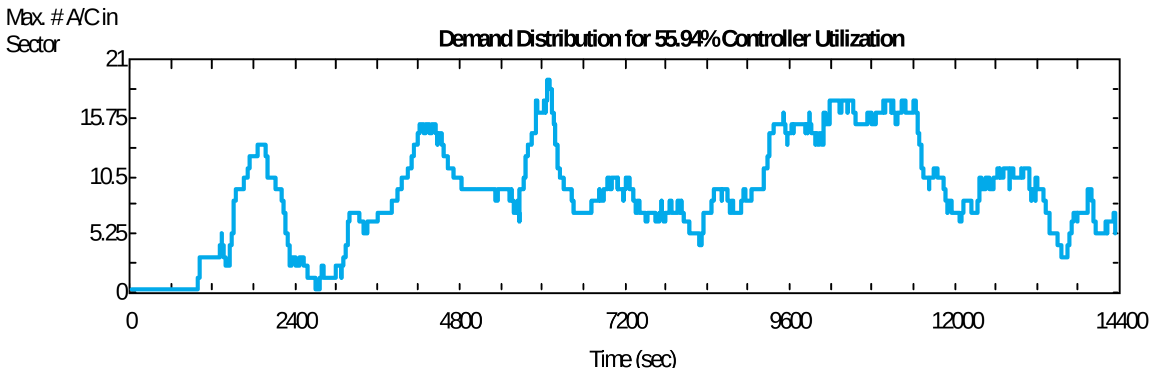

Figure 3-15. Demand Distribution for Utilization 52.82 percent

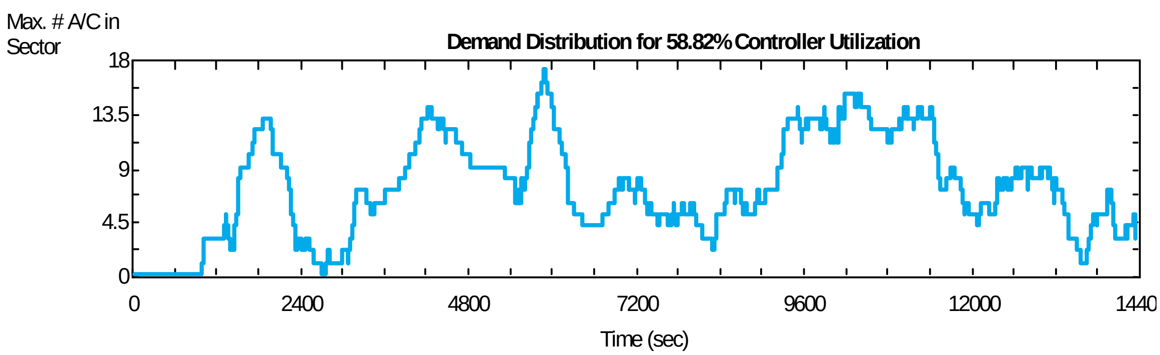

Figure 3-16. Demand Distribution for Utilization 51.65 percent

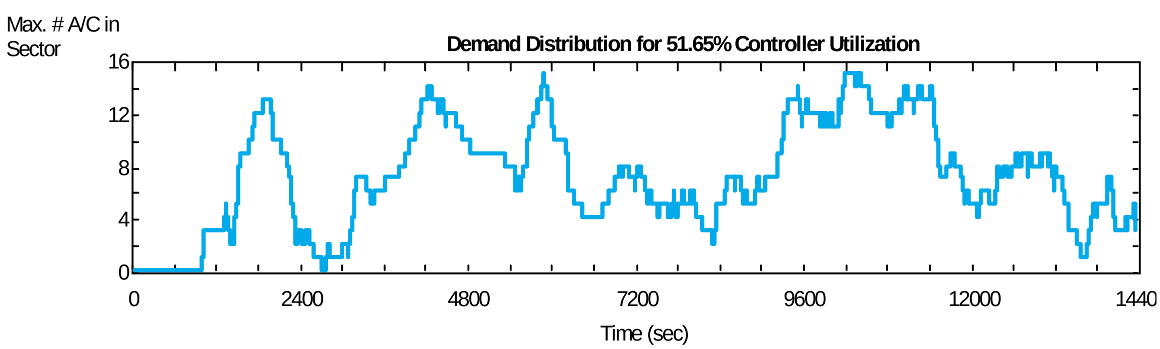

We ran the simulations for a set of event files covering a range of values of the maximum number of aircraft in the sector. This established the variation of controller utilization's with maximum number of aircraft. We repeated the simulations, using varying conflict resolution times of 5 seconds, 25 seconds, and 50 seconds per conflict. We adjusted the conflict generator so that the number of conflicts generated in a 4-hour period agreed with observations reported by Grossberg, Richards, and Robertson. [10] Figure 3-17 shows the results from the runs of the simulations. We smoothed the curves by fitting a quadratic function to them:

$$\max = c_0 + c_1(util) + c_2(util)^2 \qquad \text{[Eq. 3-1]}$$

Figure 3-17. Variation of Controller Utilization with Maximum Number of a/c in Sector

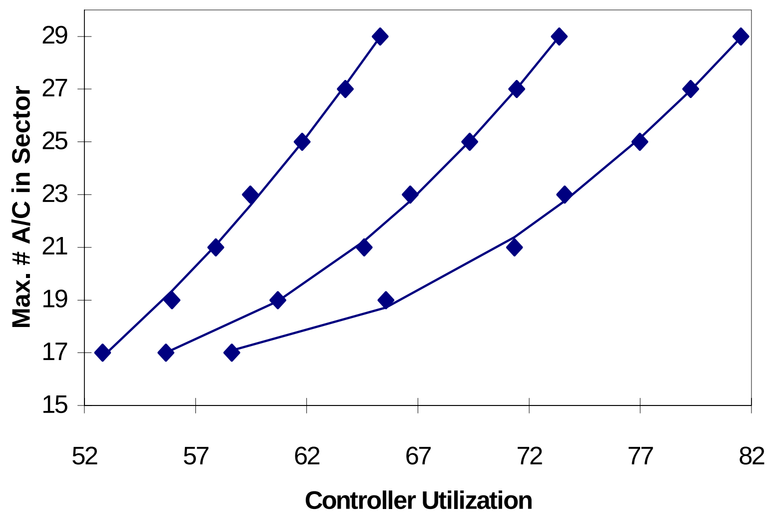

We believe a decrease in conflict resolution time from 50 seconds to 25 seconds is plausible for CPTP. When we reduced the time to resolve conflicts using a maximum number of aircraft of 17, the maximum number of aircraft able to be handled increased to 18. As the controller utilization increased, the difference between the 50-second and 25-second conflict resolution times caused a larger increase in maximum number of aircraft able to be controlled.

To conform with FAA standards [11], we use 18 as the standard maximum number of aircraft in an ARTCC or TRACON sector. While the results of Figure 3-17 might be used to justify a larger increase, to be conservative, we chose to reflect CPTP's effect as increasing the maximum number by 1 to 19.

Airspace Tool and Sector Tool

The Airspace Tool (AT) will help controllers manage traffic that passes through sectors without making transitions to or from terminal airspace. It is intended to support a new controller position, the "airspace coordinator." The airspace coordinator will have cognizance over the airspace of more than one sector, perhaps over all the sectors in a center. Using accurate forecasts of aircraft's future position, current flight plan information, and, possibly, trial planning features of the AT, the airspace coordinator will develop proposals for clearances that make efficient resolutions of conflicts, and conform closely to users' wishes. The airspace coordinator will then interact with sector controllers to implement and deliver these clearances.

The AT may be modeled with an extension of the CP/TPT model. The AT's benefits for airspace users will be principally in more efficient conflict resolutions and in clearances that are closer to the users' desired routes, in comparison with CP/TPT results.

AT's benefits to the sector controllers should exceed those of the CP/TPT because the airspace coordinator will develop even more efficient conflict resolutions than the CP/TPT, and will deliver them even more efficiently to controllers. Simulation modeling, like that of the FAM modeling reported here, may be used to develop quantitative measures of AT's benefits to sector controllers.

The Sector Tool (ST) will assist controllers managing transition airspace by developing proposals for efficient clearances. This tool will attack directly the inefficiencies noted in Chapter 2 and illustrated by the descent profile of Figure 2-23. As noted in Chapter 2, these are rife at busy terminals; eliminating them may have a substantial payoff in fuel, time, and schedule integrity.

ST's benefits to airspace users may be modeled by comparing the fuel burns and times of actual descent profiles with those of optimal descent profiles, using such a tool as the BADA/FSCM aircraft performance model. ST's benefits to air traffic managers may be modeled by simulations, for example, with FAM.

Advanced En-Route Ground Automation

This tool is intended to extend the efficiency and flexibility of ATM in en route and transitional airspace even beyond the levels provided by AT/ST. It will provide such advanced features as automatic conflict resolution, coordination among adjacent ARTCCs, and automated negotiation among ATM functions, airline AOCs, and aircrew.

Modeling Advanced En route Ground Automation (AERGA) will require extending the FAM models of ARTCC and TRACON sectors to include AOCs. FAM models aircrew workloads and functions, although we did not require this feature for the tools analyzed in this report.

Active Final Approach Spacing Tool

Describing our Active Final Approach Spacing Tool (A-FAST) model necessarily involves describing our P-FAST model. P-FAST provides controllers with advisories for landing sequence and for landing runway. As described by Davis et al. [12], the test installation of P-FAST at DFW raised the average peak arrival rate by roughly 10 percent for both IFR and VFR operations. In the baseline for that comparison, however, about three to five arrivals/hour were diverted to runways other than those in the normal set of arrival runways. Correcting for this difference in the capacity of the runways used leads to the conclusion that P-FAST caused an increase of about 13 percent in the capacity of the set of normally-used arrival runways. Reference [12] also indicates that a significant part of P-FAST's benefits were due to better balancing of the loads on separate runways.

This "thought experiment" shows that runway balancing is likely to be important at any airport with multiple runways. Suppose two independent runways are accommodating arrivals, each with a capacity of 35 arrivals/hour. Also suppose that arrival demand is 50/hour. For simplicity, let us consider steady-state operations.

If the arrival stream is evenly balanced between the two runways, each will receive 25 aircraft per hour, and, thus, will operate at a utilization ratio of 5/7. In steady state, that would cause a mean queue of 2.5 aircraft, which implies a mean delay of approximately 4.3 minutes for each arrival. Thus, with balanced runway use, the airport handles the arrival demand with delays that are significant but probably tolerable.

Now suppose there is a moderate imbalance, with the arrivals reaching the two runways in a 20-30 split. There is little delay–about 2 minutes–on the less-loaded runway but arrivals to the more heavily loaded runway will see a mean delay of more than 10 minutes. Delays of that magnitude threaten airlines' schedule integrity.

Even a slightly more serious imbalance, say, an 18-32 split, would create an intolerable 18-minute delay on the more heavily loaded runway. It is likely that flights would divert from that runway to bring delays down at least to the 10-minute level. That would imply about two diversions per hour, or a reduction in the runways' effective capacity of 4 percent.

To gain an indication of the potential effects of efficient sequencing, we considered operations for both "domestic" and "international" mixes of aircraft types. (These mixes are defined earlier in this chapter in the subsection Runway Capacities.) For the domestic mix, allowing aircraft to arrive at random gave a runway arrival rate of 32.9 per hour. Restricting the runway to just one type of aircraft gave a spread of arrival rates, ranging from 24.65 (all small) to 36.38 (all large.) Weighting each of these "one-type" arrival rates by the fraction of that type in the mix gave a weighted average arrival rate of 34.57. We take this weighted average as a crude indicator of the improvement in arrival rate that could be achieved by efficient sequencing. By this measure, efficient sequencing could increase arrival rates at domestic airports by 5 percent. Repeating the process for international airports gave an arrival rate improvement of 7 percent.

The above analyses of the effects of optimal sequencing and runway balancing suggest that sequencing and balancing together might result in about a 10 percent improvement in arrival capacity. This appears consistent with the benefits observed at DFW. Also, the analyses suggest that benefits of about that size might be expected at any airport with multiple runways, at which balancing was imperfect with present ATM methods.

Now, our airport capacity models do not include any adjustment for less-than-perfect runway balancing. In effect, they assume perfect balancing. In view of this, we model only the sequencing effect of P-FAST.

We did this by changing the Pareto parameters from those of the assigned mix to the weighted average of the one-type parameters. This led to increases of about 4 percent in departure capacity, in addition to the arrival capacity increases.

Since P-FAST is an aid to arriving traffic, that might appear to give P-FAST an unmerited effect on departures. However, DFW tests reported significant increases in departure capacity during P-FAST operations. [12] So, we are content to have our model assign some departure capacity improvements to P-FAST. In work to model effects of tools that, like A-SMA, should directly affect departure capacity, this point should be revisited so that appropriate benefits can be associated with each tool.

Now, let us consider A-FAST. A-FAST will augment the capabilities of Passive FAST with an interface that provides speed and heading advisories to the TRACON final approach controller. According to Reference 13, it also will have improved conflict-detection and resolution capabilities. Reference 13 also says that A-FAST should result in “tighter means and smaller standard deviations of in-trail separations on final approach, and shorter common approach path lengths.”

In the context of our models, we see A-FAST as giving controllers much more accurate position information for arrivals, reducing variations in approach speeds, reducing variations in approach profiles, and reducing the common approach path length.

Specifically, we model A-FAST by

- Reducing position uncertainty from 0.25 nm to 100’;
- Reducing standard deviations of approach speeds from 5 kt to 2.5 kt;
- Reducing standard deviation of wind variation from 7 kt to 5 kt (this reflects A-FAST’s reduction of variations in approach profiles); and
- Reducing common path length from 6 nm to 5 nm.

The steps in capacity from the current reference through P-FAST to A-FAST are shown in Figure 3-18, which compares the Pareto frontiers describing runway capacities in ILS Category I conditions for the three cases.

Figure 3-18. Capacity Comparisons

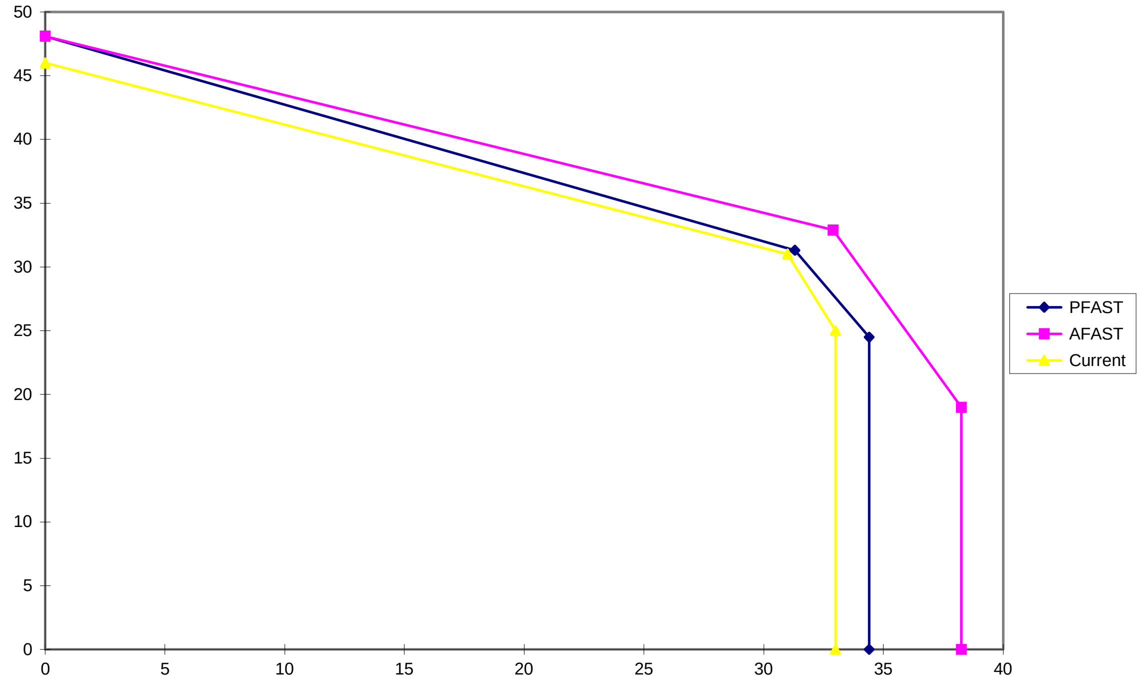

Expedite Departure Path

The performance of Expedite Departure Path (EDP) has been characterized as "decrease time-to-cruise-altitude by 15 percent." The material of Chapter 2 suggests that bringing times-to-climb for departures from busy airports to values characteristic of less-busy airports could reduce this time (specifically, the time-to-climb averaged over a day) by 3 minutes from a base of 16 minutes. That would be a decrease of 19 percent. In view of this, the 15 percent goal seems reasonable, if it is interpreted as applying only to busy airports.

EDP is to achieve its results by giving controllers suggested clearances that balance flows to departure fixes, and allow efficient climb-out paths whenever they are possible with the existing mix of arrivals and departures. Presenting controllers with suggested clearances changes their cognitive processes from doing all the work of analyzing the traffic picture and determining appropriate clearances, to reviewing the suggested clearances. This change could reduce the thinking time required for each flight. If that happened, then EDP also could increase the maximum number of aircraft that a controller can handle at one time. If so, and if also the controller's utilization is the binding constraint on the maximum number of aircraft in a particular departure TRACON, then EDP would increase the maximum number of aircraft in the departure TRACON.

Standard instrument departures from busy airports often do not have a fixed route, but, rather, instruct crews to expect vectors to one of several fixes or navigation aids. (There is just one SID for ORD, for instance, and it is of this kind.) Consequently, it seems likely that for many busy airports the controllers' utilization, rather than airspace limitations, will in fact govern the maximum number of aircraft that can be accommodated in the departure TRACON at one time.

An interview with a controller who had experience in the NYC TRACON raised a note of caution, however, about the chances for EDP to increase the number of aircraft handled at one time. The controller told us that controller teams generally developed standard operating procedures that they carried out largely mechanically, particularly during busy periods. The controller believed that this often resulted in conservative clearances. EDP operations might require controllers to do more complex tasks to issue less-conservative clearances for departures. In that case, it is not clear that the maximum number of aircraft handled could increase, even with the help provided by EDP.

A solid assessment of EDP's effects on the maximum number of aircraft simultaneously in a departure TRACON must, we believe, wait until the tool is more fully defined. Therefore, we model EDP by reducing the mean time in departure TRACONs by 3 minutes, leaving the maximum number in the sector unchanged.

Enhanced Surface Movement Advisor

Enhanced Surface Movement Advisor (E-SMA) will provide information from many sources (e.g., ARTS data, airline schedule and gate data, flight plans, ACARS data on flight status, and runway status data) to optimize the use of surface movement resources, probably by means of collaborative decision-making among surface traffic managers and airlines. Specific benefits are to include runway load balancing and managed competition for a taxiway resource.

Modeling E-SMA's benefits from runway load balancing would begin with determining how runways are assigned now. Presumably, presently each runway's load and mix are dictated by airlines' specific gates, OAG departure schedules, and a choice of taxiways made by ground controllers. With E-SMA, the runways' loads and mixes would be determined by well-informed, collaborative decisions minimizing total time from gate to wheels up in general, and giving due consideration to promoting certain flights when that is to a carrier's overall advantage.

Simulation modeling probably will be necessary to determine changes in runway loading and mixes that E-SMA would be likely to realize. With this information, LMI's runway model would capture the effects of better mixes on capacity. LMI's airport models would then determine the effects on capacity, and LMINET would capture the consequent effects on delays throughout the NAS. Another, very interesting option would be to integrate a simulation model of a specific airport or set of airports directly into LMINET.

This modeling is quite likely to be airport-specific. It should be validated by reviews with FAA controllers at each airport treated and by reviews with airline ground operations managers.

Modeling management of a scarce taxiway resource would also begin by determining how traffic reaches the resource in present operations. Presumably, now each concourse's pushback schedule is dictated by individual airlines' gates and schedules, together with decisions by the ground controller. With E-SMA, pushback schedules could be determined collaboratively, to minimize effects of congestion at the scarce resource.

These effects could be captured in LMI's airport model by introducing a queue for the scarce resource. Taxi delays can then be evaluated "before" and "after" E-SMA during operations modeled with LMINET. The LMINET calculations also will determine NAS-wide effects of installing E-SMA at specific airports.

Here, too, the situations modeled are likely to be airport-specific. The models also should be reviewed by FAA and airlines ground traffic managers.

ECONOMIC MODELS

Economic models play the key role in evaluating the relative merits of each decision support tool. They serve to translate measures of the technical effectiveness of the DSTs into monetary equivalents. We use a delta approach, comparing a baseline case (which has no DST) to a case that incorporates a particular DST. The benefits of the DSTs are assumed to be the monetary difference between the two cases. This monetary difference is also significant because it represents the marginal value of any single DST as well as relative value and importance when implementing multiple DSTs. This analysis also forms a lower estimate of an upper bound of the allowable cost of the DST. It is a lower estimate because neither qualitative economic benefits are considered nor are the second-order quantitative economic benefits.

Benefits and costs will be defined at two levels: those occurring at the air carrier/aircraft level and those occurring at the air traffic controller level. Two types of savings accrue at the air carrier/aircraft level. The first is the combined savings in flight time and fuel costs resulting from less delay. The second is the savings gained from removing part of the controller-induced ATM constraints and allowing aircraft to fly at or near optimal 4-D flight paths. At the air traffic control level, the costs and benefits are measured by the change in the number of operations handled by the sectors and the TRACONs. We made no attempt to convert these measures into equivalent staffing numbers or monetary costs; that is beyond the scope of this project.

The major output of the queuing model of the NAS is delay in aircraft-minutes. The first step in the economic analysis is to determine the cost of this delay. The approach used in this analysis was first developed by Earl Wingrove of LMI for the analysis of the Terminal Area Productivity Program. [14] Upper and lower bounds of the system-wide delay costs per block-minute of time are defined by a pessimistic and an optimistic scenario, respectively. The pessimistic estimate is based on the equipment-level direct operating cost (DOC[3]) plus an allocated share of cabin crew costs, all divided by block minutes of time. This estimate of delay costs includes fuel costs plus aircraft depreciation/amortization and rental costs. It implicitly assumes that all arrival delay occurs in the air and that some incremental capital costs are incurred during the delay period. It is the higher of the two estimates.

The optimistic scenario is based on the variable operating cost (VOC[4]), which does not include aircraft depreciation/amortization and rental costs, plus an allocated share of cabin crew costs, minus fuel costs, then all divided by block minutes of time. This estimate implicitly assumes that all arrival delay is taken on the ground and that aircraft depreciation/amortization and rental costs are not charged against the delay. It is the lower of the two estimates. Neither of these measures includes the costs to either the airlines or the flying public resulting from canceled flights. In 1995, the system-wide delay costs based on the weighted average of turboprop, short-haul jets and long-haul jets were $43.18 for the pessimistic case and $24.08 for the optimistic case.[5]

The ground hold cost is set to $42/minute. This is an approximation, based on interviews of the figures that major airlines use to price gate delays. This price consists of three major cost categories:

- Lost revenue (passengers leaving because the flight is delayed or not traveling in the future because of the delay)
- Crew and fuel costs

[3] Direct Operating Costs (DOC) are charges directly related to owning and operating the aircraft. They include flight crew costs (e.g., salaries, benefits/pensions, payroll taxes, and personnel/training expenses); fuel and oil costs (including taxes); maintenance costs (including maintenance overhead); insurance and injuries/loss/damage charges; aircraft rentals; and aircraft depreciation/amortization charges.

[4] Variable Operating Costs (VOC) are direct charges that vary as the aircraft utilization varies. They are essentially the DOC minus aircraft rentals and aircraft depreciation/amortization charges.

[5] A more complete discussion of these points is given in Chapter 6 of Reference [14].

- Downstream disruptions.

The present analysis uses the optimistic scenario as the basis for costing the system-wide delay cost. We chose the lesser estimate for a variety of reasons, but chiefly because the predicted growth in air traffic represents just one realization of many different possibilities. This particular realization does not include the competitive response of the carriers to rising delays and their associated costs. For instance, as the delays to Carrier A increase at Airport A, it is likely that Carrier A will shift some of its operations to Airport B, especially those carriers that use hub and spoke type operations. Therefore, the delays derived in this study represent a type of worst case. By using the lower estimate, we will partially compensate for that.

Much of this analysis examines delay as it occurs in various phases and modes of flight. That information is easily translated into fuel burned during that phase or mode of flight.

For operations on airports and in terminal areas, we developed representative fuel burn rates from the BADA data set. [7] The reasons for the variations in these burn rates, shown in Table 3-3, seem obvious: the burn rate for climb is of course the largest one; fuel burn for vectoring during departures ("vector out") is larger than that for vectoring during arrivals ("vector in") because the airplanes are lighter during arrivals than during departures, and so on.

Table 3-3 Fuel Burns per Flight Mode/Phase

Model	Mode/phase	Fuel burn rate (kilogram per minute)
Airport		
	Ground idle	11.01
	Taxi out	17.05
	Climb	103.64
	Vector out	50.38
	Cruise	49.46
	Vector in	33.68
	Descent	11.12
	Taxi in	15.35
Sector		
	Arrival TRACON	41.57
	Departure TRACON	11.01
	Cruise	49.46

There are three burn rates for operations in the arrival TRACON, the departure TRACON and the en route sectors. The fuel burn in the arrival TRACON is set equal to the average of fuel burn occurring in the vector in and cruise flight modes. The departure sector fuel burn is set equal to the ground idle fuel burn because we assume that delays for the departure TRACON will be taken on the ground. The en route fuel burn is set equal to the cruise fuel burn.

This economic analysis examines results that are expected to occur in 2005. As such, the benefits and costs associated with this study need to be presented in a consistent manner. Therefore, all benefits and costs need to be translated into the same dollars. We chose 1995 dollars for this reference. There are four key parameters. The initial system-wide delay costs are assumed to have a yearly increase of 1 percent per year, which represents the increase in real aircraft operating costs as the fleet ages. The real cost of fuel is assumed to increase at 0.10 percent per year and the nominal rate of inflation is set to 2 percent per year. The key parameters are inflated to the year 2005 baseline then deflated back to the equivalent 1995 dollars. The values are shown in Table 3-4.

Table 3-4. System-wide Delay Costs

Parameter	Year 2005 cost in year 2005 dollars	Year 2005 cost in 1995 dollars
Gallon of fuel	0.666	0.563
Upper bound system wide delay cost per block minute	52.632	53.685
Lower bound system wide delay cost per block minute	29.356	29.943
Ground hold delay cost per block minute		45.935

The cost of the fuel is found by multiplying the cost of the fuel in dollars per gallon times burn rate in kg per minute times the conversion factor of kilograms to gallons. This is done for each flight mode and phase. When added to the lower bound system-wide delay cost per block minute, the value used to cost block of minute of delay for this study is found. This number is shown in the last column of Table 3-5.

Table 3-5 Delay Costs

Model	Mode/phase	Delay cost in 2005 in 1995 dollars	Fuel cost in 2005 in 1995 dollars	Total delay cost per block minute
Airport				
	Ground idle	29.943	2.052	31.995
	Taxi out	29.943	3.178	33.121
	Climb	29.943	19.311	49.254
	Vector out	29.943	9.387	39.330
	Cruise	29.943	9.215	39.158
	Vector in	29.943	6.278	36.221
	Descent	29.943	2.071	32.014
	Taxi in	29.943	2.860	32.803
	Ground hold	45.935	0.00	45.935
Sector				
	Arrival TRACON	29.943	7.745	37.688
	Departure TRACON	29.943	2.052	31.995
	Cruise	29.943	9.215	39.158

Chapter 4

Some Cross-Comparable Benefits Estimates

In this chapter, we report the results of applying the instantiation of our method described in Chapter 3, to make cross-comparable estimates of the benefits of a set of three DSTs: A-FAST, CPTP, and EDP.

Decision Support Tools Considered, and Their Assumed Implementations

In this section, we describe the specific ways in which the models of A-FAST, CPTP, and EDP were used in making the benefits estimates.

Active Final Approach Spacing Tool

We applied the A-FAST model of Chapter 3 in two "builds." In the first build, we applied A-FAST only at the 10 airports at which, following a review of FAA plans, we had decided to include P-FAST (that decision, and those airports, are described in the subsection "Airport Models," in Chapter 3 page 3-3. In the second build, we applied A-FAST at all 64 LMINET airports.

Expedite Departure Path

We chose the set of TRACONs that should have EDP installations in two steps. First, we extended the analysis of ETMS data reported in Chapter 2 to determine which of the 64 network airports would benefit from the tool. As we explain in Chapter 2, ETMS data for times-to-climb show fairly sharp differences. Aircraft departing from busier airports generally require about 3 minutes' more time-to-climb than do departures from less busy airports.

The extended analysis of times-to-climb identified 20 airports in the "long-time-to-climb" category. These are ATL, CLT, CVG, DCA, DFW, EWR, IAH, JFK, LAX, LGA, MCO, MDW, MIA, MSP, ORD, PHL, PHX, SEA, SFO, and STL. Evidently, these airports are the ones most likely to benefit from EDP. The full list includes, however, several airports, such as CVG, that have significantly less traffic than the others. Eliminating these brought us to a final list of 16 airports at which EDP seems most likely to generate significant benefits: ATL, CLT, DFW, EWR, IAH, JFK, LAX, LGA, MCO, MIA, MSP, ORD, PHL, PHX, SFO, and STL. Accordingly, we implemented EDP at these airports to assess the likely benefits of that tool.

Conflict Probe/Trial Planning Tool

Demand varies greatly from sector to sector not every en route sector is likely to benefit from CPTP. As an illustration, Figure 4-1 shows peak demand from LMINET for traffic flying over the United States on the wind routes for April 8, 1996 as a function of geographic sector.

Figure 4-1. Peak ATM Demand as a Function of Geographic Sector

We chose the set of sectors in which to apply this tool by identifying the ones in which, for representative LMINET calculations, peak traffic density exceeded the FAA's standard one-controller maximum of 18. That produced the 50 sectors shown in Figure 4-2. In our assessment, we implemented CPTP at these sectors.

Figure 4-2. CPTP Sectors

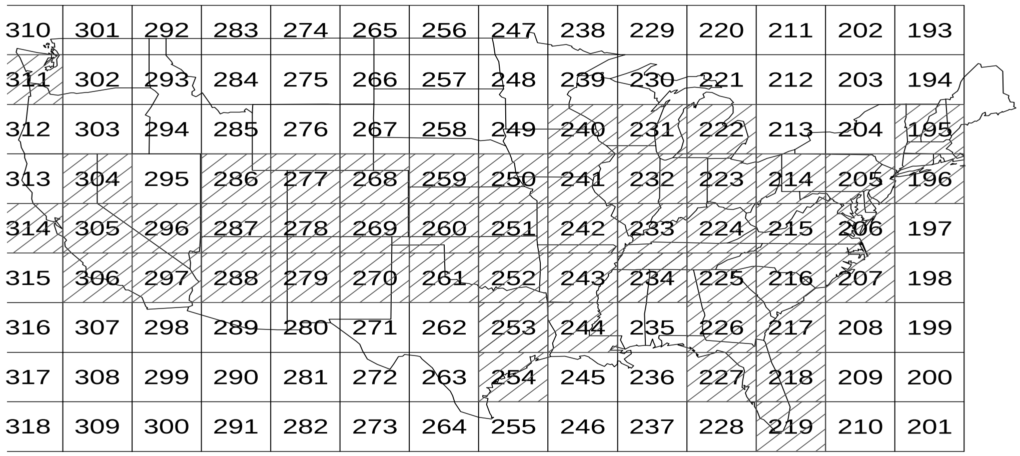

THE CROSS-COMPARABLE BENEFITS ESTIMATES

To assess the three tools, we made LMINET calculations for 30 cases. Table 4-1 lists these.

In each case, the weather inputs to the airport capacity functions were actual weather for a specific day, and the winds aloft were actual winds for that day. We considered three days: April 8,1996, June 12, 1996, and November 29, 1996.

There are six reference or baseline cases and 24 DST cases. The reference cases for 1996 were made to validate the model by comparisons with available data for that year. The reference cases for 2005 are the basis for the comparisons showing the DST's benefits.

We assigned delay costs to three causes: lack of airport arrival capacity, lack of airport departure capacity, and lack of sector capacity. Inadequate airport arrival capacity generates delays that LMINET represents in three reports. The first of these, the arrival queue, is an obvious immediate consequence of inadequate arrival capacity. Also, LMINET's model of the FAA's EDICT process generates ground holds in response to lengthy arrival queues, so LMINET's ground holds are caused by arrival delays.

Table 4-1. LMINET Cases

Case	Winds			DST			
	Good	Fair	Bad	10 A-FAST	64 A-FAST	EDP	CPTP
1		•					
2	•						
3			•				
4		•		•			
5		•			•		
6	•			•			
7	•				•		
8			•	•			
9			•		•		
10		•					•
11	•						•
12			•				•
13		•				•	
14	•					•	
15			•			•	
16		•			•		•
17	•				•		•
18			•		•		•
19		•			•	•	
20	•				•	•	
21			•		•	•	
22		•				•	•
23	•					•	•
24			•			•	•
25		•			•	•	•
26	•				•	•	•
27			•		•	•	•

Finally, if arrivals are sufficiently impeded, eventually there will not be enough airplanes for departing flights, which will cause queues for airplanes. In some circumstances, problems other than lack of arrival capacity may cause queues for planes. For example, an airport may receive a large part of its arrivals from just one other airport or small group of airports. Boston relates to the NYC airports in this way. Then, restricting departures from the "major supplier" eventually produces queues for airplanes even if the receiver's arrival capacity is adequate. These appear to be somewhat special cases, however, and we chose to associate queues for airplanes with lack of arrival capacity.

We priced delays in arrival queues at the “arrival TRACON” rate of $37.688 per minute (this rate, like all the others in this section, is developed in Chapter 3.) We costed delays in queues for planes and ground holds at the ground-hold rate, $45.935 per minute. We took delays in en route sectors to cost $39.158 per minute. Assuming that aircraft in departure queues consumed fuel at the ground idle rate, we priced these queues at $31.995 per minute.

When EDP was present, we reduced the total cost of departure capacity delays by 3 minutes at the vector-out rate of $39.330 per minute for each departure at an EDP airport.

Combining delay costs for the 3 days considered by the weights described in the Chapter 3 subsection, Modeling Annual Variations with Representative Days, these cost models generated the top-level cost results shown in Table 4-2.

Table 4-2. Estimated Annual Delay Costs in Billions of 1996 Dollars

Case	Arrival	Departure	Sector	Total	Change from baseline (percent)	Change from previous case (percent)	Absolute savings	Change in absolute savings from previous case
1996 Reference	1.517	0.278	0.118	1.913				
2005 Baseline	4.194	0.658	0.328	5.180				
2005 10 A-FAST	3.663	0.636	0.337	4.636	11		0.544	
2005 64 A-FAST	2.723	0.535	0.341	3.599	31	22	1.581	1.037
2005 64 A-FAST EDP	2.723	0.326	0.318	3.367	35	6	1.813	0.232
2005 64 A-FAST EDP CP	2.723	0.326	0.230	3.279	37	3	1.901	0.088
2005 EDP	4.194	0.449	0.305	4.948	4		0.232	
2005 EDP CP	4.194	0.449	0.220	4.863	6	2	0.317	0.085
2005 CP	4.194	0.658	0.243	5.095	2		0.085	
2005 EDP CP	4.194	0.326	0.230	4.750	8	7	0.430	0.345
2005 10 A-FAST	3.663	0.636	0.337	4.636	11		0.544	
2005 64 A-FAST	2.723	0.535	0.341	3.599	31	22	1.581	1.037
2005 64 A-FAST CP	2.723	0.535	0.253	3.511	32	2	1.669	0.088

Figure 4-3 summarizes the results.

Figure 4-3. Cross-Comparable Benefit Estimates of DSTs and Groups of DSTs

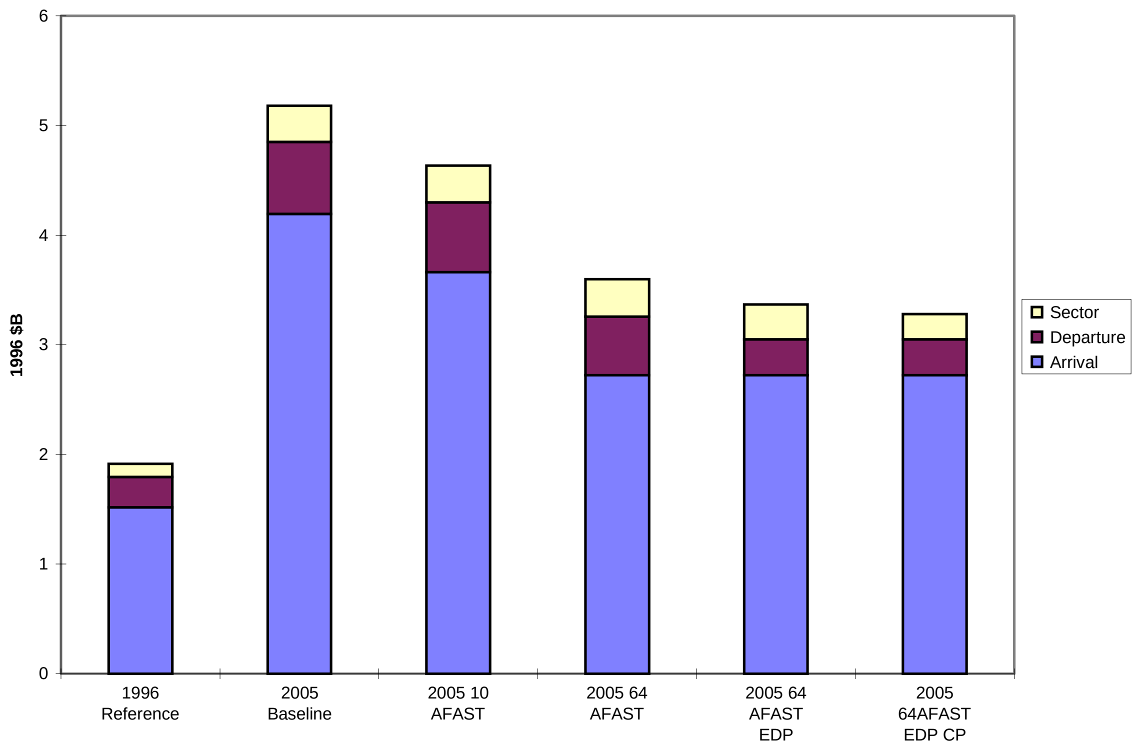

DISCUSSION

Evidently, the present instantiation of our method for assessing DSTs indicates that A-FAST is the most effective tool. This is consistent with our observations in Chapter 2, that problems associated with arrival capacity at busy airports appear to be the most prominent causes of delays and inefficiencies in the present operation of the NAS.

While the present instantiation indicates that EDP would have a smaller impact than A-FAST, the predicted annual savings of $232 million are certainly substantial. CPTP, which the present instantiation indicates would have the least effect of the three DSTs considered, nevertheless generates significant predicted savings of nearly $100 million per year.

The present results show only small interactive effects in the tools' benefits. For example, CPTP is predicted to be about 3 percent more effective when implemented with full A-FAST implementation. This is because the three tools studied affect distinct parts of the NAS. We would expect tools such as the Sector Tool, which will affect the efficiencies of both arriving and departing traffic, to show more pronounced interactions with, say, A-FAST and EDP.

It is well to remember here the warnings of Chapter 3's subsection, Some important caveats. The present instantiation predicts that delay costs would increase by a factor approaching 3, between 1996 and 2005. This might well imply delays that the airlines would find intolerable, causing them to change their operations away from the demand model of Chapter 3.

To the degree that this would happen, the total delay costs, and, perhaps, the absolute benefits of the DSTs, are overstated. Applying all three DSTs reduces the predicted increase in delay costs to "only" about 80 percent. If airlines would tolerate that kind of increase, and maintain schedules consistent with the demand model of Chapter 3, then the results of Table 4-2 are more reliable indicators of the DSTs' effects in keeping delay costs acceptable in the face of very large increases in traffic.

Chapter 5

Summary and Conclusions

To fulfill the project's chief purpose, we developed the method of Chapter 3 for generating cross-comparable estimates of the benefits of NASA-developed decision support tools. The method's key idea is to use an economic model to generate dollar-valued estimates of the reductions in delay costs and of the benefits from flying more nearly optimal flight paths, that accrue from implementing DSTs or groups of DSTs. (Dollar benefits from the tools' improvements in FAA controller productivity also could be included, although we have not done so in the present task.)

The economic model's inputs are outputs from models of the national airspace system, of air traffic management operations in ARTCCs and TRACONs, and models of aircraft performance. DST's effects are represented with adjustable features of these subordinate models.

The method for assessing DST benefits is a general one; as discussed in Chapter 3, many different models can be used for its components. We implemented the method with the particular components described in Chapter 3. With this instantiation of the method, we generated the results of Chapter 4.

From this work, we reach these conclusions:

1. The benefits of DSTs, and groups of DSTs, can be estimated consistently in dollars, with a flexible method that uses a capstone economic model to integrate outputs from subordinate models characterizing operations of the NAS.

2. System-wide delay costs are likely to increase almost threefold by 2005. This would profoundly affect airlines' operations.

3. Implementing A-FAST widely, and implementing EDP and CPTP at those terminals and ARTCC sectors where these tools have the greatest likelihood of significant payoff, would change the threefold increase to an increase of about 70 seventy percent.

4. A-FAST is the most effective tool of those we considered. Nevertheless, EDP and CPTP bring significant benefits that may well justify their implementation.

5. Implementing A-FAST at just 10 key airports would generate substantial savings.

The project's second purpose was to identify the causes of delay and inefficiency in present operations of the NAS, that NASA-developed DSTs might address. From that work, which is detailed in Chapter 2, we reach these conclusions:

1. Inadequate arrival capacity at busy terminals is the most costly present cause of delay. This is consistent with the analysis result that A-FAST, which improves arrival capacity, is the most effective of the three DSTs considered in this study.

2. Inefficient departures are a significant cause of delay.

3. Cruise routes outside terminal airspace generally are efficient.

4. Groundside congestion at certain busy airports appears to offer an opportunity for DSTs.

5. A DST that could assist controllers in making efficient insertions and extractions of aircraft into and out of holding patterns may be helpful.

Chapter 6

Recommendations for Future Work

It would seem well worthwhile to exercise the method that we have developed to assess DST benefits, to treat other tools than the three just considered. The suite of models used in the present instantiation appears adequate to treat AT/ST, so that this extension could be done fairly rapidly. Also, the present version of the method could be used to suggest optimal mixes and installation patterns for DSTs.

Treating E-SMA would require extending the models, to provide more detailed coverage of the airport groundside. To make that extension may be quite desirable, however, in view of the potential impacts of increasing groundside congestion at key terminals.

Analyzing the benefits of AERGA would also require extending the models, to cover the intended interactions among ATM, AOC, and aircrew functions.

Another profitable direction for future work would be to extend the present instantiation. One interesting and potentially quite useful direction for extensions would be to explore integrating simulation models covering parts of the NAS in considerable detail, with a queuing network. SIMMOD simulations of key airports of particular interest might provide particularly useful means of exploring the mutually interacting effects of airside and groundside tools, for example.

Another direction in which it might prove quite useful to extend the present instantiation would be to increase the number of weather-days covered, in order to generate meaningful measures of dispersion, as well as of central tendencies, of the effects of DSTs and groups of DSTs.

Finally, the method could be extended to include costs of the technologies, so that an extended economic model could be used to suggest efficient investments in DSTs.

References

[1] LMI Report NS501T1, *Technical and Economic Analysis of Air Transportation Management Issues Related to Free Flight*, Lee, D. A., et al. February 1997.

[2] Lee, D. A., *M/E3/N/N+Q and M/Ek/N/N+Q*, LMI Internal Paper, January 1997.

[3] Federal Aviation Administration, *1996 Aviation Capacity Enhancement Plan and Airport Database*, FAA's Office of System Capacity, Washington, DC 20591.

[4] LMI, *Weather and Delay Data for 10 Major Airports*, Plugge, J., October 1995.

[5] LMI, Report for Contract No. NAS2-14361, *User's Guide: Aircraft/Air Traffic Management Functional Analysis Model*, Etheridge, M. R., February 1997.

[6] Kaplan, B. J. et al. *The ASAC Flight Segment and Network Cost Models*, NASA Contractor Report 201679, Contract NAS2-14361, NASA Langley Research Center, April 1997.

[7] EUROCONTROL Experimental Centre, *User Manual for the Base of Aircraft Data (BADA) Revision 2.4*, EEC Note 5/96, February 1996.

[8] *ACSYNT User Guide and Installation Manual*, Phoenix Integration, 1872 Pratt Drive, Blacksburg, Virginia 24060.

[9] McCullers, L. A., *Flight Optimization System Release 5.91 User's Guide*, available by electronic mail. (Transfers may be arranged by telephone to L. A. McCullers at 757-864-7631, or by mail to amccul@avd00.larc.nasa.gov.)

[10] Grossberg, M., J. Richards, and A. Robertson, *Validation of Air Traffic Controller Workload Models*, September 1979.

[11] Federal Aviation Administration, FAA *Order 7210.3M*, Change 1, May 1996, U. S. Department of Transportation, Washington, D. C.

[12] Davis, J. T., D. R. Isaacson, J. E. Robinson III, W. den Braven, K. K. Lee, and B. Sanford, *Operational Test Results of the Passive Final Approach Spacing Tool*, 8th IFAC Symposium on Transportation Systems, Chania, Greece, June 1997.

[13] Ballin, M. G, R. A. Coppenbarger, D. R. Schleicher, and S. C. Johnson, *Summary Overview and Status of AATT Program Development Activities*, NASA Internal Report, draft of 29 May 1997.

[14] Lee, David A., et al. *Estimating the Effects of the Terminal Area Productivity Program*, NASA Contractor Report 201682, Contract NAS2-14361, NASA Langley Research Center, Hampton, Virginia 23681-0001.

Appendix A
Glossary of Airport Identifiers

ABQ	Albuquerque International Airport, Albuquerque, New Mexico
APA	Arapahoe Airport, Denver, Colorado
ATL	The William B. Hartsfield Atlanta International Airport, Atlanta, Georgia
AUS	Robert Mueller Airport, Austin, Texas
BDL	Bradley International Airport, Hartford, Connecticut/Springfield, Massachusetts
BNA	Nashville, Tennessee Airport
BOS	General Edward Lawrence Logan International Airport, Boston, Massachusetts
BUR	Burbank, California Airport
BWI	Baltimore-Washington International Airport
CHI	Chicago. Used to represent the two Chicago terminals, considered together
CLE	Hopkins International Airport, Cleveland, Ohio
CLT	Douglas Airport, Charlotte, North Carolina
CMH	Columbus International Airport, Columbus, Ohio
CVG	Cincinnati-Northern Kentucky Airport, Cincinnati, Ohio
DAL	Love Field, Dallas/Fort Worth, Texas
DAY	Dayton International Airport, Dayton, Ohio
DBQ	Dubuque Municipal Airport, Dubuque, Iowa
DCA	Washington National Airport, Washington, D. C.
DEN	Denver International Airport, Denver, Colorado

DFW	Dallas-Fort Worth International Airport, Dallas/Fort Worth, Texas
DTW	Detroit Metropolitan Wayne County Airport, Detroit, Michigan
EWR	Newark International Airport, Newark, Ohio
HOU	William P. Hobby Airport, Houston, Texas
HPN	Westchester County Airport, New York
IAD	Dulles International Airport, Washington, D. C.
IAH	Houston Intercontinental Airport, Houston, Texas
IND	Indianapolis International Airport, Indianapolis, Indiana
ISP	MacArthur Field, Long Island, New York
JFK	John F. Kennedy International Airport
LAS	McCarran International Airport, Las Vegas, Nevada
LAX	Los Angeles International Airport, Los Angeles, California
LGA	La Guardia Airport, New York, New York
LGB	Daugherty Field, Long Beach, California
MCI	Kansas City International Airport, Kansas City, Missouri
MCO	Orlando International Airport, Orlando, Florida
MDW	Midway Airport, Chicago, Illinois
MEM	Memphis International Airport, Memphis, Tennessee
MIA	Miami International Airport, Miami, Florida
MKE	General Mitchell Field, Milwaukee, Wisconsin
MSP	Minneapolis-Saint Paul International Airport, Minneapolis-Saint Paul, Minnesota
MSY	New Orleans International Airport, New Orleans, Louisiana
NYC	New York City (used to represent the three New York terminals considered as a group.)
OAK	Oakland International Airport, Oakland, California

OMA	Eppley Airport, Omaha, Nebraska
ONT	Ontario International Airport, Ontario, California
ORD	Chicago O Hare International Airport
PBI	Palm Beach International Airport, Palm Beach, Florida
PDX	Portland International Airport, Portland, Oregon
RDU	Raleigh/Durham International Airport, Raleigh-Durham, North Carolina
RNO	Reno-Tahoe Airport, Reno, Nevada
SAN	Lindbergh Field, San Diego, California
SAT	San Antonio International Airport, San Antonio, Texas
SDF	Standiford Field, Louisville, Kentucky
SEA	Seattle-Tacoma International Airport, Seattle, Washington
SFO	San Francisco International Airport, San Francisco, California
SJC	San Jose Airport, San Jose, California
SLC	Salt Lake City International Airport, Salt Lake City, Utah
SMF	Sacramento Metropolitan Airport, Sacramento, California
SNA	John Wayne International Airport, Orange County, California
STL	Lambert Field, Saint Louis, Missouri
SYR	Hancock Field, Syracuse, New York
TEB	Teterboro Airport, Teterboro, New Jersey
TPA	Tampa International Airport, Tampa-St. Petersburg, Florida
YVR	Vancouver International Airport, Vancouver, British Columbia, Canada

Appendix B
Glossary of Acronyms

AATT	Advanced Air Transportation Technologies
ACARS	ARINC Communications Addressing and Reporting System
ACE	Aviation Capacity Enhancement
ACSYNT	Aircraft Synthesis Model
AERGA	Advanced Enroute Ground Automation
AFAST	Active Final Approach Spacing Tool
AND	Approximate Network Delays
AOC	Air Operations Center
ARINC	Aeronautical Radio, Incorporated
ARTCC	Air Route Traffic Control Center
ARTS	Automated Radar Terminal System
ASAC	Aviation Systems Analysis Capability
ASD	Aircraft Situation Display
ASMA	Advanced Surface Movement Advisor
ASQP	Airline Service and Quality Performance
AT	Airspace Tool
AT/ST	Airspace Tool/Sector Tool
ATC	Air Traffic Control
ATM	Air Traffic Management
BADA	Base of Aircraft Data
CONUS	Contiguous United States

CP	Conflict Probe
CP/TP	Conflict Probe/Trial Planning Tool
CPTPT	Conflict Probe/Trial Planning Tool
CTAS	Center-TRACON Automation System
DOC	Direct Operating Cost
DPAT	Detailed Policy Analysis Tool
DST	Daylight Saving Time
EDCT	Estimated Departure Clearance Time
EDP	Expedite Departure Path
EEC	European Economic Community
ELP	El Paso International Airport, El Paso, Texas
ESMA	Enhanced Surface Movement Advisor
ETMS	Enhanced Traffic Management System
EURO-CONTROL	European Organization for the Safety of Air Navigation
FAA	Federal Aeronautics Administration
FAM	Functional Analysis Model
FAST	Final Approach Spacing Tool
FSCM	Flight Segment Cost Model
FL	Flight Level
FLL	Fort Lauderdale International Airport, Fort Lauderdale, Florida
FLOPS	Flight Optimization System
FOD	Fort Dodge Airport, Fort Dodge, Iowa
FSCM	Flight Segment Cost Model
GSO	Greensboro/High Point/Winston-Salem Airport, North Carolina

HCT	Hayes Center VORTAC
IBM	International Business Machines
IFAC	International Federation of Automatic Control
IFR	Instrument Flight Rules
ILS	Instrument Landing System
IMC	Instrument Meteorological Conditions
LMINET	A queuing network model of the U. S. national airspace system
MIT	Miles-in-Trail; also Massachusetts Institute of Technology
NAS	National Airspace System
NASPAC	National Airspace System Performance Analysis Capability
NASSIM	National Airspace Simulation
NCEP	National Center for Environmental Prediction
OAG	Official Airline Guide
OASIS	National Climatic Data Center's On-Line Access and Service Information System
PAMS	Performance Analysis Monitoring System
PFAST	Passive Final Approach Spacing Tool
PMAC	Performance Monitoring Analysis Capability
RAMS	Reorganized ATC Mathematical Simulator
RISC	Reduced Instruction Set Computer
SID	Standard Instrument Departure
SIMMOD	FAA-developed Airport and Airspace Simulation Model
STAR	Standard Terminal Arrival Route
TAAM	Total Airport and Airspace Modeler
TAP	Terminal Area Productivity

TOC	Top of Climb
TOD	Top of Descent
TRACON	Terminal Radar Approach Control
UAL	United Air Lines
VFR	Visual Flight Rules
VHF	Very High Frequency
VMC	Visual Meteorological Conditions
VOC	Variable Operating Cost
VOR	VHF Omni Range
VORTAC	VOR/Tactical Air Navigation
ZDV	Symbol for the Denver Air Route Traffic Control Center

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY ***(Leave blank)***

2. REPORT DATE
July 1998

3. REPORT TYPE AND DATES COVERED
Contractor Report

4. TITLE AND SUBTITLE
A Method for Making Cross-Comparable Estimates of the Benefits of Decision Support Technologies for Air Traffic Management

5. FUNDING NUMBERS
C NAS2-14361

WU 538-16-11-01

6. AUTHOR(S)
David Lee, Dou Long, Mel Etheridge, Joana Plugge, Jesse Johnson and Peter Kostiuk

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Logisitics Management Institute
2000 Corporation Ridge
McLean, Virginia 22102-7805

8. PERFORMING ORGANIZATION REPORT NUMBER
NS710S1

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Langley Research Center
Hampton, VA 23681-0001

10. SPONSORING / MONITORING AGENCY REPORT NUMBER
NASA/CR-1998-208455

11. SUPPLEMENTARY NOTES
Langley Technical Monitor: Robert E. Yackovetsky
Final Report

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Unclassified - Unlimited

Subject Category 01
Availability: NASA CASI (301) 621-0390
Distribution: Nonstandard

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 words)***

We present a general method for making cross comparable estimates of the benefits of NASA-developed decision support technologies for air traffic management and we apply a specific implementation, or "instantiation" of the method to estimate benefits of three decision support tools, Active Final Approach Spacing Tool (AFAST), Expedite Departure Path (EDP), and Conflict Probe and Trial Planning Tool (CPTP). The report also reviews data about the present operation of the national airspace system (NAS) to identify opportunities for DST's to reduce delays and inefficiencies.

14. SUBJECT TERMS
Air Traffic Management (ATM)
Advanced Air Transportation Technologies (AATT)
Air Transporation
EDP
CPTP
AFAST

15. NUMBER OF PAGES
92

16. PRICE CODE
A05

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	Unlimited

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102